

Randomness is Unpredictability

Antony Eagle

Published in *British Journal for the Philosophy of Science*, **56** (2005), 749–90.

Abstract

The concept of randomness has been unjustly neglected in recent philosophical literature, and when philosophers have thought about it, they have usually acquiesced in views about the concept that are fundamentally flawed. After indicating the ways in which these accounts are flawed, I propose that randomness is to be understood as a special case of the epistemic concept of the unpredictability of a process. This proposal arguably captures the intuitive desiderata for the concept of randomness; at least it should suggest that the commonly accepted accounts cannot be the whole story and more philosophical attention needs to be paid.

1 Randomness in Science

- 1.1 Random Systems
- 1.2 Random Behaviour
- 1.3 Random Sampling
- 1.4 Caprice, Arbitrariness, and Noise

2 Concepts of Randomness

- 2.1 Von Mises/Church/Martin-Löf Randomness
- 2.2 KCS-randomness

3 Randomness is Unpredictability: Preliminaries

- 3.1 Process and Product Randomness
- 3.2 Randomness is Indeterminism?

4 Predictability

- 4.1 Epistemic Constraints on Prediction
- 4.2 Computational Constraints on Prediction
- 4.3 Pragmatic Constraints on Prediction
- 4.4 Prediction Defined

5 Unpredictability

6 Randomness is Unpredictability

6.1 Clarification of the Definition of Randomness

6.2 Randomness and Probability

6.3 Subjectivity and Context-Sensitivity of Randomness

7 Evaluating the Analysis

[R]andomness [...] is going to be a concept which is relative to our body of knowledge, which will somehow reflect what we know and what we don't know.

HENRY E. KYBURG, JR. ([1974], p. 217)

Phenomena that we cannot predict must be judged random.

PATRICK SUPPES ([1984], p. 32)

The concept of randomness has been sadly neglected in the recent philosophical literature. As with any topic of philosophical dispute, it would be foolish to conclude from this neglect that the truth about randomness has been established. Quite the contrary: the views about randomness in which philosophers currently acquiesce are fundamentally mistaken about the nature of the concept. Moreover, since randomness plays a significant role in the foundations of a number of scientific theories and methodologies, the consequences of this mistaken view are potentially quite serious. After I briefly outline the scientific roles of randomness, I will survey the false views that currently monopolise philosophical thinking about randomness. I then make my own positive proposal, not merely as a contribution to the correct understanding of the concept, but also hopefully prompting a renewal of philosophical attention to randomness.

The view I defend, that randomness is unpredictability, is not entirely without precedent in the philosophical literature. As can be seen from the epigraphs I quoted at the beginning of this paper, the connection between the two concepts has made an appearance before.¹ These quotations are no more than suggestive however: these

authors were aware that there is some kind of intuitive link, but made no efforts to give any rigorous development of either concept in order that we might see how and why randomness and prediction are so closely related. Indeed, the Suppes quotation is quite misleading: he adopts exactly the pernicious hypothesis I discuss below (§3.2), and takes determinism to characterise predictability—so that what he means by his apparently friendly quotation is exactly the mistaken view I oppose! Correspondingly, the third objective I have in this paper is to give a plausible and defensible characterisation of the concept of predictability, in order that we might give philosophical substance and content to this intuition that randomness and predictability have something or other to do with one another.²

1 Randomness in Science

The concept of randomness occurs in a number of different scientific contexts. If we are to have any hope of giving a philosophical concept of randomness that is adequate to the scientific uses, we must pay some attention to the varied guises in which randomness comes.

All of the following examples are in some sense derivative from the most central and crucial appearance of randomness in science—randomness as a prerequisite for the applicability of probabilistic theories. Von Mises was well aware of the centrality of this role; he made randomness part of his *definition* of probability. This association of randomness with von Mises' hypothetical frequentism has unfortunately meant that interest in randomness has declined with the fortunes of that interpretation of probability. As I mentioned, this decline was hastened by the widespread belief that randomness can be explained merely as indeterminism. Both of these factors have led to the untimely neglect of randomness as a centrally important concept for understanding a number of issues, among them being the ontological force of probabilistic theories, the criteria and grounds for acceptance of theories, and how we might evaluate the strength of various proposals concerning statistical inference. Especially when one considers the manifest inadequacies of ontic accounts of randomness when dealing with these issues (§2), the neglect of the concept of randomness seems to have left a significant gap in the foundations of probability. We should, however, be wary of associating worries about randomness too closely with

issues in the foundations of probability—those are only one aspect of the varied scientifically important uses of the concept. By paying attention to the use of the concept, hopefully we can begin to construct an adequate account that genuinely plays the role required by science.

1.1 Random Systems

Many dynamical processes are modelled probabilistically. These are processes which are modelled by probabilistic state transitions.³ Paradigm examples include the way that present and future states of the weather are related, state transitions in thermodynamics and between macroscopic partitions of classical statistical mechanical systems, and many kinds of probabilistic modelling. Examples from ‘chaos theory’ have been particularly prominent recently (Smith, [1998]).

For example, in ecohydrology (Rodriguez-Iturbe, [2000]), the key concept is the soil water balance at a point within the rooting depth of local plants. The differential equations governing the dynamics of this water balance relate the rates of rainfall, infiltration (depending on soil porosity and past soil moisture content), evapotranspiration and leakage (Laio *et al.*, [2001]; Rodriguez-Iturbe *et al.*, [1999]). The occurrence and amount of rainfall are random inputs.⁴ The details are interesting, but for our purposes the point to remember is that the randomness of the rainfall input is important in explaining the robust structure of the dynamics of soil moisture. Particular predictions of particular soil moisture based on particular volumes of rainfall are not nearly as important for this project as understanding the responses of soil types to a wide range of rainfall regimes.

The robust probabilistic structures that emerge from low-level random phenomena are crucial to the task of explaining and predicting how such systems evolve over time and what consequences their structure has for the systems that depend on soil moisture, for example, plant communities.⁵

Similar dynamical models of other aspects of the natural world, including convection currents in the atmosphere, the movement of leaves in the wind, and the complexities of human behaviour, are also successfully modelled as processes driven by random inputs. But the simplest examples are humble gaming devices like coins

and dice. Such processes are random if anything is: the sequence of outcomes of heads and tails of a tossed coin exhibits disorder, and our best models of the behaviour of such phenomena are very simple probabilistic models.

At the other extreme is the appearance of randomness in the outcomes of systems of our most fundamental physics: quantum mechanical systems. Almost all of the interpretations of quantum mechanics must confront the randomness of experimental outcomes with respect to macroscopic variables of interest; many account for such randomness by positing a fundamental random process. For instance, collapse theories propose a fundamental stochastic collapse of the wave function onto a particular determinate measurement state, whether mysteriously induced by an observer (Wigner, [1961]), or as part of a global indeterministic dynamics (Ghirardi *et al.*, [1986]; Bell, [1987a]). Even no-collapse theories have to claim that the random outcomes are not reducible to hidden variables.⁶

1.2 Random Behaviour

The most basic model that population genetics provides for calculating the distribution of genetic traits in an offspring generation from the distribution of such traits in the parent generation is the Hardy-Weinberg Law (Hartl, [2000], p. 26–9).⁷ This law idealises many aspects of reproduction; one crucial assumption is that mating between members of the parent generation is random. That is, whether mating occurs between arbitrarily selected members of the parent population does not depend on the presence or absence of the genetic traits in question in those members. Human mating, of course, is not random with respect to many genetic traits: the presence of particular height or skin colour, for example, does influence whether two human individuals will mate. But even in humans mating is random with respect to some traits: blood group, for example. In some organisms, for instance some corals and fish, spawning is genuinely random: the parent population gathers in one location and each individual simply ejects their sperm or eggs into the ocean where they are left to collide and fertilise; which two individuals end up mating is a product of the random mixing of the ocean currents. The randomness of mating is a pre-requisite for the application of the simple dynamics; there is no explicit presence of a random state transition, but such behaviour is presupposed in the application of the theory.

Despite its many idealisations, the Hardy-Weinberg principle is explanatory of the dynamics of genetic traits in a population. Complicating the law by making its assumptions more realistic only serves to indicate how various unusual features of actual population dynamics can be deftly explained as a disturbance of the basic underlying dynamics encoded in the law. As we have seen, for some populations, the assumptions are not even idealised. Each mating event is random, but nevertheless the overall distribution of mating is determined by the statistical features of the population as a whole.

Another nice example of random behaviour occurs in game theory. In many games where players have only incomplete information about each other, a randomising mixed strategy dominates any pure strategy (Suppes, [1984], pp. 210–2). Another application of the concept of randomness is to agents involved in the evolution of conventions (Skyrms, [1996], pp. 75–6). For example, in the convention of stopping at lights which have two colours but no guidance as to the intended meaning of each, or in the reading of certain kinds of external indicators in a game of chicken (hawk-dove), the idea is that the players in the game can look to an external source, perceived as random, and take that as providing a way of breaking a symmetry and escaping a non-optimal mixed equilibrium in favour of what Aumann calls a *correlated equilibrium*. In this case, as in many others, the epistemic aspects of randomness are most important for its role in scientific explanations.

1.3 Random Sampling

In many statistical contexts, experimenters have to select a representative sample of a population. This is obviously important in cases where the statistical properties of the whole population are of most interest. It is also important when constructing other experiments, for instance in clinical or agricultural trials, where the patients or fields selected should be independent of the treatment given and representative of the population from which they came with respect to treatment efficacy. The key assumption that classical statistics makes in these cases is that the sample is *random* (Fisher, [1935]).⁸ The idea here, again, is that we should expect no correlation between the properties whose distribution the test is designed to uncover and the properties that decide whether or not a particular individual should be tested.

In Fisher's famous thought experiment, we suppose a woman claims to be able to taste whether milk was added to the empty cup or to the tea. We wish to test her discriminatory powers; we present her with eight cups of tea, exactly four of which had the milk added first. The outcome of a trial of this experiment is a judgement by the woman of which cups of tea had milk. The experimenter must strive to avoid correlation between the order in which the cups are presented, and whatever internal algorithm the woman uses to decide which cups to classify as milk-first. That is, he must randomise the cup order. If her internal algorithm is actually correlated with the presence of milk-first, the randomising should only rule out those cases where it is not, namely, those cases where she is faking it.

An important feature of this case is that it is important that the cup selection be random to the woman, but not to the experimenters. The experimenters want a certain kind of patternlessness in the ordering of the cups, a kind of disorder that is *designed* to disturb accidental correlations (Dembski, [1991]). The experimenters also wish themselves to know in what order the cups are coming; the experiment would be uninterpretable without such knowledge. Intuitively, this order would not be random for the experimenters: they know which cup comes next, and they know the recipe by which they computed in which order the cups should come.

1.4 Caprice, Arbitrariness, and Noise

John Earman has argued that classical Newtonian mechanics is indeterministic, on the basis of a very special kind of case (Earman, [1986], pp. 33–39). Because Newtonian physics imposes no upper bound on the velocity of a point particle, it is nomologically possible in Newtonian mechanics to have a particle whose velocity is finite but unbounded, which appears at spatial infinity at some time t (this is the temporal inverse of an unboundedly accelerating particle that limits to an infinite velocity in a finite time). Prior to t that particle had not been present in the universe; hence the prior state does not determine the future state, since such a 'space invader' particle is possible. Of course, such a space invader is completely unexpected—it is plausible, I think, to regard such an occurrence as completely and utterly random and arbitrary. Randomness in this case does not capture some potentially explanatory aspect of

some process or phenomenon, but rather serves to mark our recognition of complete capriciousness in the event.

More earthly examples are not hard to find. Shannon noted that when modelling signal transmission systems, it is inappropriate to think that the only relevant factors are the information transmitted and the encoding of that information (Shannon & Weaver, [1949]). There are physical factors that can corrupt the physical representation of that data (say, stray interference with an electrical or radio signal). It is not appropriate or feasible to explicitly incorporate such disturbances in the model, especially since they serve a purely negative role and cannot be controlled for, only accommodated. These models therefore include a *random noise* factor: random alterations of the signal with a certain probability distribution. All the models we have mentioned include noise as a confounding factor, and it is a very general technique for simulating the pattern of disturbances even in deterministic systems with no other probabilistic aspect. The randomness of the noise is crucial: if it were not random, it could be explicitly addressed and controlled for. As it stands, noise in signalling systems is addressed by complex error-checking protocols, which, if they work, rely crucially on the random and unsystematic distribution of errors. A further example is provided by the concept of random mutation in classical evolutionary theory. It may be that, from a biochemical perspective, the alterations in DNA produced by imperfect copying are deterministic. Nevertheless, these mutations are random with respect to the genes they alter, and hence the differential fitness they convey.⁹

2 Concepts of Randomness

If we are to understand what randomness is, we must begin with the scientifically acceptable uses of the concept. These form a rough picture of the intuitions that scientists marshal when describing a phenomenon as random; our task is to systematise these intuitions as best we can into a rigorous philosophical analysis of this intuitive conception.

Consider some of the competing demands on an analysis of randomness that may be prompted by our examples.

1. **Statistical Testing** We need a concept of randomness adequate for use in random sampling and randomised experiments. In particular, we need to be able to produce random sequences on demand, and ascertain whether a given sequence is random.
2. **Finite Randomness** The concept of randomness must apply to the single event, as in Earman's example or a single instance of random mating. It must at least apply to finite phenomena.
3. **Explanation and Confirmation** Attributions of randomness must be able to be explanatorily effective, indicating why certain systems exhibit the kinds of behaviour they do; to this end, the hypothesis that a system is random must be amenable to incremental empirical confirmation or disconfirmation.
4. **Determinism** The existence of random processes must be compatible with determinism; else we cannot explain the use of randomness to describe processes in population genetics or chaotic dynamics.

Confronted with this variety of uses of randomness to describe such varied phenomena, one may be tempted to despair: 'Indeed, it seems highly doubtful that there is anything like a unique notion of randomness there to be explicated.' (Howson & Urbach, [1993], p. 324). Even if one recognises that these demands are merely suggested by the examples and may not survive careful scrutiny, this temptation may grow stronger when one considers how previous explications of randomness deal with the cases we described above. This we shall now do with the two most prominent past attempts to define randomness: the place selection/statistical test conception and the complexity conception of randomness. Both do poorly in meeting our criteria; poorly enough that if a better account were to be proposed, we should reject them.

2.1 Von Mises/Church/Martin-Löf Randomness

Definition 1 (von Mises-Randomness). An infinite sequence \mathbf{S} of outcomes of types A_1, K, A_n , is *vM-random* iff (i) every outcome type A_i has a well-defined relative

frequency $relf_i^S$ in S ; and (ii) for every infinite subsequence S^* chosen by an *admissible place selection*, the relative frequency of A_i remains the same as in the larger sequence: $relf_i^S = relf_i^{S^*}$.

Immediately, the definition only applies to infinite sequences, and so fails condition (2) of finiteness.

What is an admissible place selection? Von Mises himself says:

[T]he question whether or not a certain member of the original sequence belongs to the selected partial sequence should be settled *independently of the result* of the observation, i.e. before anything is known about the result. ([1957], p. 25).

The intuition is that, if we pick out subsequences independently of the contents of the elements we pick (by paying attention only to their indices, for example), and each of those has the same limit relative frequencies of outcomes, then the sequence is random. If we could pick out a biased subsequence, that would indicate that some set of indices had a greater than chance probability of being occupied by some particular outcome; the intuition is that such an occurrence would not be consistent with randomness.

Church ([1940]), attempting to make von Mises' remarks precise, proposed that admissible place selections are recursive functions that decide whether an element s_i is to be included in a subsequence on input of the index number i and the initial segment of the sequence up to s_{i-1} . For example, 'select only the odd numbered elements', and 'select any element that comes after the subsequence 010' are both admissible place selections. An immediate corollary is that no random sequence can be recursively computable: else there would be a recursive function that would choose all and only 1s from the initial sequence, namely, the function which generates the sequence itself. But if a random sequence cannot be effectively generated, we cannot produce random sequences for use in statistical testing. Neither can we effectively test, for some given sequence, whether it is random. For such a test would involve exhaustively checking *all* recursive place selections to see whether they produce a deviant subsequence, and this is not decidable in any finite time (though for any non-random sequence, at some finite time the checking machine will halt with a 'no')

answer). If random sequences are neither producible or discernible, they are useless for statistical testing purposes, failing the first demand. This point may be made more striking by noting that actual statistical testing only ever involves finite sequences; and no finite sequence can be vM-random at all.

Furthermore, it is perfectly possible that some genuinely vM-random infinite sequence has an arbitrarily biased initial segment, even to the point where all the outcomes of the sequence that actually occur during the life of the universe are 1s. A theorem of Ville ([1939]) establishes a stronger result: given any countable set of place selections $\{\varphi_i\}$, there is some infinite sequence \mathbf{S} such that the limit frequency of 1s in any subsequence of $\mathbf{S}^* = \varphi_j(\mathbf{S})$ selected by some place selection is one half, despite the fact that for every finite initial segment of the sequence, the frequency of 1s is greater than or equal to one half (van Lambalgen, [1987], pp. 730–1, 745–8). That is, any initial segment of this sequence is not random with respect to the whole sequence or any infinite selected subsequence. There seems to be no empirical constraint that could lead us to postulate that such a sequence is genuinely vM-random. Indeed, since any finite sequence is recursively computable, no finite segment will ever provide enough evidence to justify claiming that the actual sequence of outcomes of which it is a part is random. That our evidence is at best finite means that the claim that an actual sequence is vM-random is empirically underdetermined, and deserving of an arbitrarily low credence because any finite sequence is better explained by some other hypothesis (e.g. that it is produced by some pseudo-random function). vM-randomness is a profligate hypothesis that we cannot be justified in adopting. Hence it can play no role in explanations of random phenomena in finite cases, where more empirically tractable hypotheses will do far better.

One possible exception is in those cases where we have a rigorous demonstration that the behaviour in question cannot be generated by a deterministic system—in that case, the system may be genuinely vM-random. Even granting the existence of such demonstrations, note that in this case we have made essential appeal to a fact about the random process that produces the sequence, and we have strictly gone beyond the content of the evidence sequence in making that appeal. Here we have simply abandoned the quest to explain deterministic randomness. Random sequences may well exist for infinite strings of quantum mechanical measurement

outcomes, but we don't think that random phenomena are confined to indeterministic phenomena alone: vM-randomness fails demand (4).

Partly in response to these kinds of worries, a final modification of Von Mises' suggestion was made by Martin-Löf ([1966], [1970], [1969]). His idea is that biased sequences are possible but unlikely: non-random sequences, including the types of sequences considered in Ville's theorem, form a set of measure zero in the set of all infinite binary sequences. Martin-Löf's idea is that truly random sequences satisfy all the probability one properties of a certain canonical kind: recursive sequential significance tests—this means (roughly) that a sequence is random with respect to some hypothesis H_p about probability p of some outcome in that sequence if it does not provide grounds for rejecting H_p at arbitrarily small levels of significance.¹⁰ Van Lambalgen ([1987]) shows that Martin-Löf (ML)-random sequences are, with probability 1, vM-random sequences also—almost all strictly increasing sets of integers (Wald place selections) select infinite subsequences of a random sequence that preserve relative frequencies.

Finally, as Dembski ([1991]) points out, for the purposes of statistical testing, 'Randomness, properly to be randomness, must leave nothing to chance.' (p. 75) This is the idea that in constructing statistical tests and random number generators, the first thing to be considered is the kinds of patterns that one wants the random object to avoid instantiating. Then one considers the kinds of objects that can be constructed to avoid matching these tests. In this case, take the statistical tests you don't want your sequence to fail, and make sure that the sequence is random with respect to these patterns. Arbitrary segments of ML-random sequences cannot satisfy this requirement, since they must leave up to chance exactly which entities come to constitute the random selection.

2.2 KCS-randomness

One aspect of random sequences we have tangentially touched on is that random sequences are intuitively complex and disordered. Random mating is disorderly at the level of individuals; random rainfall inputs are complex to describe. The other main historical candidate for an analysis of randomness, suggested by the work of

Kolmogorov, Chaitin and Solomonov (KCS), begins with the idea that randomness is the (algorithmic) complexity of a sequence.¹¹

The *complexity* of a sequence is defined in terms of effective production of that sequence.

Definition 2 (Complexity). The complexity $\mathbf{K}_T(\mathbf{S})$ of sequence \mathbf{S} is the length of the shortest programme C of some Turing machine \mathbf{T} which produces \mathbf{S} as output, when given as input the length of \mathbf{S} . ($\mathbf{K}_T(\mathbf{S})$ is set to ∞ if there does not exist a C that produces \mathbf{S}).

This definition is machine-dependent; some Turing machines are able to more compactly encode some sequences. Kolmogorov showed that there exist universal Turing machines \mathbf{U} such that for any sequence \mathbf{S} ,

$$(1) \quad \forall \mathbf{T} \exists c_T (\mathbf{K}_U(\mathbf{S}) \leq \mathbf{K}_T(\mathbf{S}) + c_T),$$

where the constant c_T doesn't depend on the length of the sequence, and hence can be made arbitrarily small as the length of the sequence increases. Such machines are known as *asymptotically optimal* machines. If we let the complexity of a sequence be defined relative to such a machine, we get a relatively machine-independent characterisation of complexity.¹² The upper bound on complexity of a sequence of length l is approximately l —we can always resort to hard-coding the sequence plus an instruction to print it.

Definition 3 (KCS-Randomness). A sequence \mathbf{S} is *KCS-random* if its complexity is approximately its length: $\mathbf{K}(\mathbf{S}) \approx l(\mathbf{S})$.¹³

One natural way to apply this definition to physical processes is to regard the sequence to be generated as a string of successive outcomes of some such process. In dynamical systems, this would naturally be generated by examining *trajectories* in the system: sequences that list the successive cells (of some partition of the state space) that are traversed by a system over time. KCS-randomness is thus primarily a property of trajectories. This notion turns out to be able to be connected with a

number of other mathematical concepts that measure some aspects of randomness in the context of dynamical systems.¹⁴

This definition fares markedly better with respect to some of our demands than vM-randomness. Firstly, there are finite sequences that are classified as KCS-random. For each l , there are 2^l binary sequences of length l . But the non-KCS-random sequences amongst them are all generated by programmes of less than length $l - k$, for some k ; hence there will be at most 2^{l-k} programmes which generate non-KCS-random sequences. But the fraction $2^{l-k}/2^l = 1/2^k$; so the proportion of non-KCS-random sequences within all sequences of length l (for all l) decreases exponentially with the degree of compressibility demanded. Even for very modest compression in large sequences (say, $k = 20$, $l = 1000$) less than 1 in a million sequences will be non-KCS-random. It should, I think, trouble us that, by the same reasoning, longer sequences are more KCS-random. This means that single element sequences are not KCS-random, and so the single events they represent are not KCS-random either.¹⁵

It should also disturb us that biased sequences are less KCS-random than unbiased sequences (Earman, [1986], pp. 143–5). A sequence of tosses of a biased coin (e.g. $\Pr(H) > 0.5$) can be expected to have more frequent runs of consecutive 1s than an unbiased sequence; the biased sequence will be more compressible. A single 1 interrupting a long sequence of 0s is even less KCS-random. But in each of these cases, intuitively, the distribution of 1s in the sequence can be as random as desired, to the point of satisfying all the statistical significance tests for their probability value. This is important because stochastic processes occur with arbitrary underlying probability distributions, and randomness needs to apply to all of them: intuitively, random mating would not be less random were the distribution over genotypes non-uniform.

What about statistical testing? Here, again, there are no effective computational tests for KCS-randomness, nor any way of effectively producing a KCS-random sequence.¹⁶ This prevents KCS-random sequences being effectively useful in random sampling and randomisation. Furthermore, the lack of an effective test renders the hypothesis of KCS-randomness of some sequence relatively immune to confirmation or disconfirmation.

One suggestion is that perhaps we were mistaken in thinking that KCS complexity is an analysis of randomness; perhaps, as Earman ([1986]) suggests, it actually is an analysis of *disorder* in a sequence, irrespective of the provenance of that sequence. Be that as it may, the problems above seem to disqualify KCS-randomness from being a good analysis of randomness. (Though random phenomena typically exhibit disorderly behaviour, and this may explain how these concepts became linked, this connection is neither necessary or sufficient.)

3 Randomness is Unpredictability: Preliminaries

Perhaps the foregoing survey of mathematical concepts of randomness has convinced you that no rigorously clarified concept can meet every demand on the concept of randomness that our scientific intuitions place on it. Adopting a best candidate theory of content (Lewis, [1984]), one may be drawn to the conclusion that no concept perfectly fills the role delineated by our four demands, and one may then settle on (for example) KCS-randomness as the best partial filler of the randomness role.

Of course this conclusion only follows if there is no better filler of the role. I think there is. My hypothesis is that scientific randomness is best analysed as a certain kind of *unpredictability*. I think this proposal can satisfy each of the demands that emerge from our quick survey of scientific applications of randomness. Before I can state my analysis in full, some preliminaries need to be addressed.

3.1 Process and Product Randomness

The mathematical accounts of randomness we addressed do not, on the surface, make any claims about scientific randomness. Rather, these accounts invite us to infer, from the randomness of some sequence, that the process underlying that sequence was random (or that an event produced by that process and part of that sequence was random). Our demands were all constraints on random processes, requiring that they might be used to randomise experiments, to account for random behaviour, and that they might underlie stochastic processes and be compatible with determinism. Our true concern, therefore, is with process randomness, not product randomness (Earman, [1986], pp. 137–8). Our survey showed that the inference from product to

process randomness failed: the class of processes that possess vM-random or KCS-random outcome sequences fails to satisfy the intuitive constraints on the class of random processes.¹⁷

Typically, appeals are made at this point to theorems which show that ‘almost all’ random processes produce random outcome sequences, and vice versa (Frigg, [2004], p. 431). These appeals are beside the point. Firstly, the theorems depend on quite specific mathematical details of the models of the systems in question, and these details do not generalise to all the circumstances in which randomness is found, giving such theorems very limited applicability. Secondly, even where these theorems can be established, there remains a logical gap between process randomness and product randomness, some random processes exhibit highly ordered outcomes. Such a possibility surely contradicts any claim that product randomness and process randomness are ‘extensionally equivalent’ (Frigg, [2004], p. 431).

What is true is that product randomness is a defeasible incentive to inquire into the physical basis of the outcome sequence, and it provides at least a prima facie reason to think that a process is random. Indeed, this presumptive inference may explain much of the intuitive pull exercised by the von Mises and KCS accounts of randomness. For, insofar as these accounts do capture typical features of the outputs of random processes, they can appear to give an analysis of randomness. But this presumptive inference can be defeated; and even the evidential status of random products is less important than it seems—on my account, far less stringent tests than von Mises or KCS can be applied that genuinely do pick out the random processes.

3.2 Randomness is Indeterminism?

The comparative neglect of the concept of randomness by philosophers is in large part due, I think, to the pervasive belief in the pernicious hypothesis that a physical process is random just when that process is indeterministic. Hellman, while concurring with our conclusion that no mathematical definition of random sequence can adequately capture physical randomness, claims that ‘physical randomness’ is ‘roughly interchangeable with “indeterministic”’ (Hellman, [1978], p. 83).

Indeterminism here means that the complete and correct scientific theory of the process is indeterministic. A scientific theory we take to be a class of models (van

Fraassen, [1989], ch. 9). An individual model will be a particular history of the states that a system traverses (a specification of the properties and changes in properties of the physical system over time): call such a history a *trajectory* of the system. The class of all possible trajectories is the scientific theory. Two types of constraints govern the trajectories: the dynamical laws (like Newton's laws of motion) and the boundary conditions (like the Hamiltonian of a classical system restricts a given history to a certain allowable energy surface) govern which states can be accessed from which other states, while the laws of coexistence and boundary conditions determine which properties can be combined to form an allowable state (for instance, the idea gas law $PV = nRT$ constrains which combinations of pressure and volume can coexist in a state). This model of a scientific theory is supposed to be very general: the states can be those of the phase space of classical statistical mechanics, or the states of soil moisture, or of a particular genetic distribution in a population, while the dynamics can include any mappings between states.¹⁸

Definition 4 (Earman-Montague Determinism). A scientific theory is *deterministic* iff any two trajectories in models of that system which overlap at one point overlap at every point. A theory is *indeterministic* iff it is not deterministic; equivalently, if two systems can be in the same state at one time and evolve into distinct states. A system is (in)deterministic iff the theory which completely and correctly describes it is (in)deterministic. (Earman, [1986]; Montague, [1974])

Is it plausible that the catalogue of random phenomena we began with can be simply unified by the assumption that randomness is indeterminism? It seems not. Many of the phenomena we enumerated do not seem to depend for their randomness on the fact that the world in which they are instantiated is one where quantum indeterminism is the correct theory of the microphysical realm. One can certainly imagine that Newton was right. In Newtonian possible worlds, the kinds of random phenomena that chaotic dynamics gives rise to are perfectly physically possible; so too with random mating, which depends on a high-level probabilistic hypothesis about the structure of mating interactions, not low-level indeterminism.¹⁹ Our definition of indeterminism made no mention of the concept of probability; an adequate understanding of randomness, on the other hand, must show how

randomness and probability are related—hence indeterminism cannot be randomness. Moreover, we must at least allow for the possibility that quantum mechanics will turn out to be deterministic, as on the Bohm theory (Bell, [1987b]). Finally, it seems wrong to say that coin tossing is indeterministic, or that creatures engage in indeterministic mating: it would turn out to be something of a philosophical embarrassment if the only analysis our profession could provide made these claims correct.

One response of behalf of the pernicious hypothesis is that, while classical physics is deterministic, it is nevertheless, on occasion, a useful idealisation to pretend that a given process is indeterministic, and hence random.²⁰ I think that this response confuses the content of concepts deployed within a theory, like the concept of randomness, with the external factors that contribute to the adoption of a theory, such as that theory being adequate for the task at hand, and therefore being a useful idealisation. Classical statistical mechanics does not say that it is a useful idealisation that gas motion is random; the theory is an idealisation that says gas motion is random, simpliciter. Here, I attempt to give a characterisation of randomness that is uniform across all theories, regardless of whether those theories are deployed as idealisations or as perfectly accurate descriptions.

We must also be careful to explain why the hypothesis that randomness is indeterminism seems plausible to the extent that it does. I think that the historical connection of determinism with prediction in the Laplacean vision can explain the intuitive pull of the idea that randomness is objective indeterminism. I believe that a historical mistake still governs our thinking in this area, for when increasing conceptual sophistication enabled us to tease apart the concepts of determinism and predictability, randomness remained connected to determinism, rather than with its rightful partner, predictability. It is to the concept of predictability that we now turn.

4 Predictability

Laplace's vision is that determinism is idealised predictability:

[A]n intelligence which could comprehend all the forces by which nature is

animated and the respective situation of all the [things which] compose it—an intelligence sufficiently vast to submit these data to analysis—it would embrace in the same formula the movements of the greatest bodies in the universe and those of the lightest atom; for it, nothing would be uncertain and the future, as well as the past, would be present to its eyes. ([1951], p. 4)

Definition 5 (Laplacean Determinism). A system is *Laplacean deterministic* iff it would be possible for an epistemic agent who knew precisely the instantaneous state and could analyse the dynamics of that system to predict with certainty the entire precise trajectory of the system.

A Laplacean deterministic system is where the epistemic features of some ideal agent cohere perfectly with the ontological features of that world. Given that there are worlds where prediction and determinism mesh in this way, it is easy to think that prediction and determinism are closely related concepts.²¹

There are two main ways to make the features of this idealised epistemic agent more realistic that would undermine this close connection. The first way is to try and make the epistemic capacities of the agent to ascertain the instantaneous state more realistic. The second way is to make the computational and analytic capacities of the agent more realistic. Weakening the epistemic abilities of the ideal agent allows us to clearly see the separation of predictability and determinism.²²

4.1 Epistemic Constraints on Prediction

The first kind of constraint to note concerns our ability to precisely ascertain the instantaneous state of a system. At best, we can establish that the system was in a relatively small region of the state space, over a relatively short interval of time.

There are several reasons for this. Most importantly, we humans are limited in our epistemic capabilities. Our measurement apparatus is not capable of arbitrary discrimination between different states, and is typically only able to distinguish properties that correspond to quite coarse partitions of the state space. In the case of classical statistical mechanics of an ideal gas in a box, the standard partition of the

state space is into regions that are macroscopically distinguishable by means of standard mechanical and thermodynamic properties: pressure, temperature, volume. We are simply not capable of distinguishing states that can differ by arbitrarily little: one slight shift of position in one particle in a mole of gas. In such cases, with even one macrostate compatible with more than one indistinguishable microstate, predictability for us and determinism do not match; our epistemic situation is typically worse than this.²³

There is an ‘in principle’ restriction too. Measurement involves interactions: a system must be disturbed, ever so slightly, in order for it to affect the system that is our measurement device. We are forced to meddle and manipulate the natural world in ways that render uncertain the precise state of the system. This has two consequences. Firstly, measurement alters the state of the system, meaning we are never able to know the precise pre-measurement state (Bishop, [2003], §5). This is even more pressing if we consider the limitations that quantum mechanics places on simultaneous measurement of complementary quantities. Secondly, measurement introduces errors into the specification of the state. Repetition does only so much to counter these errors; physical magnitudes are always accompanied by their experimental margin of error.

It would be a grave error to think that the in principle limitations are the more significant restrictions on predictions. On the contrary: prediction is an activity that arose primarily in the context of agency, where having reasonable expectations about the future is essential for rational action. Creatures who were not goal directed would have no use for predictions. As such, an adequate account of predictability must make reference to the actual abilities of the epistemic agents who are deploying the theories to make predictions. An account of prediction which neglected these pragmatic constraints would thereby leave out why the concept of prediction is important or interesting at all (Schurz, [1995], §6).

A nice example of the consequences of imprecise specification of initial conditions is furnished by the phenomenon from chaotic dynamics known as *sensitive dependence on initial conditions*, or ‘error inflation’ (Smith, [1998], pp. 15, 167–8). Consider some small bundle of initial states S , and some state $s_0 \in S$. Then, for some systems,

$$(2) \forall \varepsilon > 0 \exists \delta > 0 \exists s_0^* \in S \exists t > 0 \left(|s_0 - s_0^*| < \delta \wedge |s_t - s_t^*| > \varepsilon \right)$$

That is, for some bundle of state space points that are within some arbitrary distance in the state space, there are at least two states whose subsequent trajectories diverge by at least after some time t . In fact, for typically *chaotic* systems, all neighbouring trajectories within the bundle of states diverge exponentially fast. Predictability fails; knowledge of initial macrostates, no matter how fine grained, can always leave us in a position where the trajectories traversing the microstates that compose that initial macrostate each end up in a completely different macrostate, giving us no decisive prediction.

How well can we accommodate this behaviour? It turns out then that predictability in such cases is exponentially expensive in initial data; to predict even one more stage in the time evolution of the system demands an exponential increase in the accuracy of the initial state specification. Given limits on the accuracy of such a specification, our ability to predict will run out in a very short time for lots of systems of very moderate complexity of description, even if we have the computational abilities. However (and this will be important in the sequel) we can predict global statistical behaviour of a bundle of trajectories. This is typically because our theory yields probabilities of state transitions from one macrostate into another.²⁴ This combination of global structure and local instability is an important conceptual ingredient in randomness (Smith, [1998], ch. 4). Bishop ([2003]) makes the plausible claim that any error in initial measurement will eventually yield errors in prediction, but exponential error inflation is a particularly spectacular example.

4.2 Computational Constraints on Prediction

There may also be constraints imposed by our inability to track the evolution of a system along its trajectory. Humphreys' ([1978]) purported counterexamples to the thesis that randomness is indeterminism relied on the following possibility: that the total history of a system may supervene on a single state, hence the system is deterministic, while no computable sequence of states is isomorphic to that history. Given the very plausible hypothesis that human predictors have at best the computation capacities of Turing machines, this means that some state evolutions are

not computable by predictors like us. This is especially pronounced when the dynamical equations governing of the system are not integrable and do not admit of a closed-form solution (Stone, [1989]). Predictions of future states when the dynamics are based on open-form solutions are subject to ever-increasing complexity as the time scale of the prediction increases.

There is a sense in which all deterministic systems are computable: each system does effectively produce its own output sequence. If we are able (*per impossibile*) to arbitrarily control the initial conditions, then we could use the system itself as an ‘analogue computer’ that would simulate its own future behaviour. This, it seems to me, would be prediction by cheating. What we demand of a prediction is the making of some reasonable, theoretically-informed judgement about the unknown behaviour of a system—not remembering how it behaved in the past. (Similarly, predicting by consulting a reliable oracle is not genuine prediction either.) I propose that, for our purposes, we set prediction by cheating aside as irrelevant.

An important issue for computation of predictions is the internal discrete representation of continuous physical magnitudes; this significant problem is completely bypassed by analogue computation (Earman, [1986], ch. VI). This approach also neglects more mundane restrictions on computations: our finite lifespan, resources, memory and patience!

4.3 Pragmatic Constraints on Prediction

There are also constraints placed on prediction by the structure of the theory yielding the predictions. Consider thermodynamics. This theory gives perfectly adequate dynamical constraints on macroscopic state conditions. But it does not suffice to predict a state that specifies the precise momentum and position of each particle; those details are ‘invisible’ to the thermodynamic state. Some features of the state are thus unpredictable because they are not fixed by the theory’s description of the state.

This is only of importance because, on occasion, this is a desirable feature of theory construction. A theory of population genetics might simply plug in the proviso that mating happens unpredictably, where this is to be taken as saying that, for the purposes of the explanatory and predictive tasks at hand, it can be effectively treated as such. It is more perspicuous not to attempt to explain this higher-order stochastic

phenomenon in terms of lower level theories. This is part of a general point about the explanatory significance of higher-level theories, but it has particular force for unpredictability. Some theories don't repay the effort required to make predictions using them, even if those theories could, in principle, predict with certainty. Other theories are more simple and effective because various deterministic phenomena are treated as absolutely unpredictable. A random aspect of the process is perhaps to be seen as a qualitative factor in explanation of some quite different phenomenon; or as an ancillary feature not of central importance to the theory; or it might simply be proposed as a central irreducible explanatory hypothesis, whose legitimacy derives from the fruitfulness of assuming it. Given that explanation and prediction are tasks performed by agents with certain cognitive and practical goals in hand (van Fraassen, [1980]), the utility of some particular theory for such tasks will be a matter of pragmatic qualities of the theory.

4.4 Prediction Defined

Given these various constraints, I will now give a general characterisation of the predictability of a process.

Definition 6 (Prediction). A *prediction function* $\Psi_{P,T}(M, t)$ takes as input the current state M of a system described by a theory T as discerned by a predictor P , and an elapsed time parameter, and yields a temporally-indexed probability distribution Pr_t over the space of possible states of the system. A *prediction* is a specific use of some prediction function by some predictor on some initial state and elapsed time, who then adopts Pr_t as their posterior credence function (conditional on the evidence and the theory). (If the elapsed time is negative, the use is a *retrodiction*.)

Let us unpack this a little. Consider a particular system that has been ascertained to be in some state M at some time. The states are supposed to be distinguished by the epistemic capacities of the predictors, so that in classical mechanics, for example, the states in question will be *macrostates*, individuated by differences in observable parameters such as temperature or pressure. A prediction is an attempt to establish what the probability is that the system will be in some other

state after some time t has elapsed.²⁵ The way such a question is answered, on my view, is by deploying a function of a kind whose most general form is a prediction function. The agent P who wishes to make the prediction has some epistemic and computational capabilities; these delimit the fine-grainedness of the partition of which M is a member, and the class of possible functions. The theory T gives the basic ingredients for the prediction function, establishing the physical relations between states of the theory accepted by the agent. These are *contextual* features that are fixed by the surroundings in which the prediction is made: the epistemic and computational limitations of the predictor and the theory being utilised are presuppositions of the making of a prediction (Stalnaker, [1984]). These contextual features fix a set of prediction functions that are available to potential predictors in that context. The actual prediction, however, is the updating of credences by the predictor who conditions on observed evidence and accepted theory, which jointly dictate the prediction functions that are available to the predictor.

The notion of an available prediction function may need some explanation. Clearly, the agent who updates by simply picking some future event and giving it credence 1 is updating his beliefs in future outcomes in a way that meets the definition of a prediction function. Nevertheless, this prediction function is (most likely) inconsistent with the theory the agent takes to most accurately describe the situation he is concerned to predict, unless that agent adopts a very idiosyncratic theory. As such, it is accepted theory and current evidence which are to be taken as basic; these fix some prediction functions as reasonable for the agents who believe those theories and have observed that evidence, and it is those reasonable prediction functions that are available to the agent in the sense I have discussed here. Availability must be a normative notion; it cannot be, for example, that a prediction function is available if an agent could update their credences in accordance with its dictates; it must also be reasonable for the agent to update in that way, given their other beliefs.²⁶

Graham Priest suggested to me that the set of prediction functions be all recursive functions on the initial data, just so as to make the set of available predictions the same for all agents. But I don't think we need react quite so drastically, especially since to assume the availability of these functions is simply to reject some of the plausible computational limitations on human predictions.

This conception of prediction has its roots in consideration of classical statistical mechanics, but the use of thermodynamic macrostates as a paradigm for the input state M may skew the analysis with respect to other theories.²⁷ The input state M must include all the information we currently possess concerning the system whose behaviour is to be predicted. This might include the past history of the system, for example when we use trends in the stockmarket as input to our predictive economic models. It must also include some aspects of the microstate of the system, as in quantum mechanics, where the uniform initial distribution over phase space in classical statistical mechanics is unavailable, so all probabilities of macroscopic outcomes are state dependent. Sometimes we must also include relevant knowledge or assumptions about other potentially interacting systems. This holds not only in cases where we assume that a system is for all practical purposes closed or isolated, but also in special relativity, where we can only predict future events if we impose boundary conditions on regions spacelike separated from us (and hence outside our epistemic access), for example that those regions are more or less like our past light cone. So the input state must be broader than just the current observations of the system, and it must include all the ingredients necessary, whatever those might be, to fix on a posterior probability function.

The relation of the dynamical equations of the theory to the available prediction functions is an important issue. The aim of a predictive theory is to yield useful predictions by means of a modified dynamics that is not too false to the underlying dynamics. For some theories, the precise states will be ascertainable and the dynamical equations solvable; the prediction functions in this case will just be the dynamical equations used in the theory, and the probability distribution over final states will be concentrated on a point in the deterministic case, or given by the basic probabilistic rule in the indeterministic case (say, Born's rule in elementary quantum theory). Other cases are more complicated. In classical statistical mechanics, we have to consider how the entire family of trajectories that intersect M (i.e. overlap the microstates s that constitute M) behave under the dynamical laws, and whether tractable functions that approximate this behaviour can be found. For instance, the very simple prediction function for ergodic statistical mechanical systems is that the probability of finding a system in some state M after sufficient time has elapsed is the proportion of the phase space that M occupies. This requires a great many

assumptions and simplifications, ergodicity prominent among them, and each theory will have different requirements. The general constraints seem to be those laid down in the preceding subsections, but no more detailed universal recipe for producing prediction functions can be given. In any case, the particular form of prediction functions is a matter for physical theory; the logical properties of such a function are those I have specified above.

Of course, whether any function that meets these formal requirements is a useful or good prediction function is another matter. A given prediction function can yield a distribution that gives probability one to the whole state space, but no information about probabilities over any more fine grained partition. Such a function, while perfectly accurate, is pragmatically useless and should be excluded by contextual factors. In particular, I presume that the predictor wishes to have the most precise partition of states that is compatible with accurate prediction. But the tradeoff between accuracy and fine-grainedness will depend on the situation in hand.

The ultimate goal, of course, is that the probability distribution given by the prediction function will serve as normative for the credences of the agents making the prediction (van Fraassen, [1989], p. 198). The probabilities are matched with the credence by means of a *probability coordination* rule, of which the Principal Principle is the best known example (Lewis, [1980]). This is essential in explaining how predictions give rise to action, and is one important reason why the outcomes of a prediction must be probabilistic. Another is that we can easily convert a probability distribution over states into an *expectation* value for the random variable that represents the unknown state of the system. Prediction can then be described as yielding expectation values for some system given an estimation of the current values that characterise the system, which enables a large body of statistical methodology to come to bear on the use and role of predictions.²⁸

5 Unpredictability

With a characterisation of predictability in hand, we are in a position to characterise some of the ways that predictability can fail. Importantly, since we have separated predictability from determinism, it turns out that being indeterministic is one way, but not the only way, in which a phenomenon can fail to be predictable.

Definition 7 (Unpredictability). An event E (at some temporal distance t) is *unpredictable* for a predictor P iff P 's posterior credence in E after conditioning on current evidence and the best prediction function available to P is not 1—that is, if the prediction function yields a posterior probability distribution that doesn't assign probability 1 to E .²⁹

There is some worry that this definition is too inclusive—after all, there are many future events that are intuitively predictable and yet we are not certain that they will occur. This worry can be assuaged by attending to the following two considerations. Firstly, this definition captures the idea that an event is not perfectly predictable. If the available well-confirmed prediction function allows us to considerably raise our posterior credence in the event, we might well be willing to credit it with significant predictive powers, even though it does not convey certainty on the event. This only indicates that between perfect predictability, and the kind of unpredictability we shall call randomness (below, §6), there are greater or lesser degrees of unpredictability. Often, in everyday circumstances, we are willing to collapse some of these finer distinctions: we are willing, for example, to make little distinction between certainty and very high non-unity credences. (This is at least partially because the structure of rational preference tends to obscure these slight differences which make no practical difference to the courses of action we adopt to achieve our preferred outcomes.) It is therefore readily understood that common use of the concept of unpredictability should diverge from the letter, but I suggest not the spirit, of the definition given above. Secondly, we must recognise that when we are prepared to use a theory to predict some event, and yet reserve our assent from full certainty in the predictions made, what we express by that is some degree of uncertainty regarding the theory. Our belief in and acceptance of theories is a complicated business, and we frequently make use of and accept theories that we do not believe to be true. Some of what I have to say here about pragmatic factors involved in prediction reflects the complexities of this matter. But regardless of our final opinion on acceptance and use of theories, it remains true that our conditional credences concerning events, conditional on the truth of those theories, capture the important credential states as far as predictability is concerned. So, many events are

predictable according to the definition above, because conditional credence in the events is 1, conditional on the simple theories we use to predict them. But we nevertheless refrain from full certainty because we are not certain of the simple theory. The point is that prediction as I've defined it concerns what our credences would be if we discharged the condition on those credences, by coming to believe the theory with certainty; and this obviously simplifies the actual nature of our epistemic relationship with the theories we accept.

An illustration of the definition in action is afforded by the case of indeterminism, the strongest form of unpredictability. If the correct theory of some system is indeterministic, then we can imagine an epistemic agent of perfect computational and discriminatory abilities, for whom the contextually salient partition on state space individuates single states, and who believes the correct theory. An event is unpredictable for such an agent just in case knowledge of the present state does not concentrate posterior credence only upon states containing the event. If the theory is genuinely indeterministic there exist lawful future evolutions of the system from the current state to each of incompatible future states S and S^* . If there is any event true in S but not in S^* , that event will be unpredictable. Indeed, if an indeterministic theory countenances any events that are not instantiated everywhere in the state space, then those events will be unpredictable.

It is important to note that predictability, while relative to a predictor, is a theoretical property of an event. It is the available prediction functions for some given theory that determine the predictions that can be made from the perspective of that theory. It is the epistemic and computational features of predictors that fix the appropriate theories for them to accept—namely, predictors accept theories which partition the state space at the right level of resolution to fit their epistemic capacities, and provide prediction functions which are well-matched to their computational abilities. In other words, the level of resolution and the allowed computational expenditure are parameters of predictability, and there will be different characteristic or typical parameters for creatures of different kinds, in different epistemic communities. This situation provides another perspective on the continued appeal of the thesis that randomness is indeterminism. Theories which describe unpredictable phenomena, on this account, *treat those phenomena as indeterministic*. The way that the theory represents some situation s is the same as the theory represents some

distinct situation s^* , but the way the theory represents the future evolutions of those states s_t and s_t^* are distinct, so that within the theory we have duplicate situations evolving to distinct situations.

It is easy to see how the features that separate prediction from determinism also lead to failures of predictability. The limited capacities of epistemic agents to detect differences between fundamental detailed states, and hence their limitation to relatively coarse-grained partitions over the state space, lead to the possibility of diverging trajectories from a single observed coarse state even in deterministic systems. Then there will exist events that do not get probability one and are hence unpredictable. Note that one and the same *type* of event can be predicted at one temporal distance, and unpredictable at another, if the diverging trajectories require some extended interval of time to diverge from each other.

If the agent does not possess the computational capacities to utilise the most accurate prediction functions, they may be forced to rely on simplified or approximate methods. If these techniques do lead to predictions of particular events with certainty, then either (contra the assumption) the prediction function is not a simplification or approximation at all, or the predictions will be incorrect, and the prediction functions should be rejected. To avoid rejecting prediction functions that make incorrect but close predictions, those functions should be made compatible with the observed outcomes by explicitly considering the margins of error on the approximate predictions. Then the outputs of such functions can include the actual outcome, as well as various small deviations from actuality—they avoid conclusive falsification by predicting approximately which state will result. If such approximate predictions can include at least a pair of mutually exclusive events, then we have unpredictability with respect to those events.

Finally, if the agent accepts a theory for pragmatic reasons, then that may induce a certain kind of failure of predictability, because the agent has restricted the range of available prediction functions to those that are provided by the theory subject to the agent's epistemic and computational limitations. An agent who uses thermodynamics as his predictive theory in a world where classical statistical mechanics is the correct story of the microphysics thereby limits her ability to predict outcomes with perfect accuracy (since there are thermodynamically indistinguishable states that can evolve into thermodynamically distinguishable outcomes, if those

initial states are statistical-mechanically distinguishable). Theories also make certain partitions of the real state space salient to predictors (the so-called *level of description* that the theory operates at), and this can lead to failures of predictability in much the same way as epistemic restrictions can (even though the agents might have other, pragmatic, reasons for adopting those partitions as salient—for instance, the explanatory value of robust macroscopic accounts).

6 Randomness is Unpredictability

We are now in a position to discuss my proposed analysis. The views suggested by Suppes and Kyburg in the epigraphs to this paper provide some support for this proposal—philosophical intuition obviously acknowledges some epistemic constraints on legitimate judgements of randomness. I think that these epistemic features, derived from pragmatic and objective constraints on human knowledge, exhaust the concept of randomness.

As I discussed earlier, some events which satisfy my definition of unpredictability are only mildly unpredictable. For instance, if the events are distinguished in a fine-grained way, and the prediction concentrates the posterior probability over only two of those events, then we may have a very precise and accurate prediction, even if not perfect. These failures of prediction do not, intuitively, produce randomness. So what kind of unpredictability do I think randomness is?

The following definition captures my proposal: *randomness is maximal unpredictability*.

Definition 8 (Randomness). An event E is random for a predictor P using theory T iff E is *maximally unpredictable*. An event E is maximally unpredictable for P and T iff the posterior probability of E , yielded by the prediction functions that T makes available, conditional on current evidence, is equal to the prior probability of E . This also means that P 's posterior credence in E , conditional on theory and current evidence (the current state of the system), must be equal to P 's prior credence in E conditional only on theory.

We may call a *process* random, by extension, if each of the outcomes of the process are random. So rainfall inputs constitute a random process because the timing and magnitude of each rainfall event is random.³⁰ That is, since the outcomes of a process $\{E_1, K, E_n\}$ partition the event space, the posterior probability distribution (conditional on theory and evidence) is identical to the prior probability distribution.³¹

This definition and its extension immediately yields another, very illuminating, way to characterise randomness: a random event is probabilistically *independent* of the current and past states of the system, given the probabilities supported by the theory (when those current and past states are in line with the coarse-graining of the event space appropriate for the epistemic and pragmatic features of the predictor). The characteristic random events, on this construal, are the successive tosses of a coin: independent trials, identically distributed because the theory which governs each trial is the same, and the current state is irrelevant to the next or subsequent trials—a so-called Bernoulli process. But the idea of randomness as probabilistic independence is of far wider application than just to these types of cases, since any useful prediction method aims to uncover a significant correlation between future outcomes and present evidence, which would give probabilistic dependence between outcomes and input states. This connection between unpredictability and probabilistic independence is in large part what allows our analysis to give a satisfactory account of the statistical properties of random phenomena. I regard it as a significant argument in favour of my account that it can explain this close connection.

However, there are a number of processes for which a strict probabilistic independence assumption fails. For example, though over long time scales the weather is quite unpredictable, from day to day the weather is more stable: a fine day is more likely to be followed by another fine day. Weather is not best modelled by a Bernoulli process, but rather by a Markov process, that is, one where the probability of an outcome on a trial is explicitly dependent on the current state. Indeed, probably most natural processes are not composed of a sequence of independent events. Independence of events in a system is likely only to show itself over timescales where sensitive dependence on initial conditions and simplified dynamics have time to compound errors to the point where nothing whatsoever can be reliably inferred from the present case to some quite distant future event.³² The use of ‘random’ to describe those processes which may display some short term predictability is quite in order,

once we recognise the further contextual parameter of the temporal distance between input state and event (or random variable) to be predicted, and that for quite reasonable timescales these processes can become unpredictable. (This also helps us decide *not* to classify as random those processes which are unpredictable in the limit as t grows arbitrarily, but which are remarkably regular and predictable at the timescales of human experimenters.) That the commonsense notion of randomness includes such partially unpredictable processes is a *prima facie* reason to take unpredictability, not independence, to be the fundamental notion—though nothing should obscure the fact that probabilistic independence is the most significant aspect of unpredictability for our purposes.³³

It is a central presupposition of my view that we can make robust statistical predictions concerning any process, random or not.³⁴ One of the hallmarks of random processes is that these are the best reliable predictions we can make, since the expectations of the variables whose values describe the characteristics of the event are well defined even while the details of the particular outcomes are obscure prior to their occurrence. This is crucial for the many scientific applications of randomness: random selections are unpredictable with respect to the exact composition of a sample (the event), but the overall distribution of properties over the individuals in that sample is supposed to be representative of the frequencies in the population as a whole. In random mating, the details of each mating pair are not predictable, but the overall rates of mating between parents of like genotype is governed by the frequency of that genotype in the population.³⁵

I wish to emphasise again the role of theories. An event is random, just if it is unpredictable, that is, if the best theoretical representation of that event relative to a given predictor leaves the probability of that event unchanged when conditioned on the current state and the laws of the theory. We should give a naturalised account of the best theory relative to a predictor: that theory should be the one that maximises fit between the epistemic and computational capacities of the predictors and the demands on those capacities made by the theory, where those capacities are perfectly objective features of the predictors. An event is random, then, just in case these objective features of the agents in question render that event unpredictable.³⁶ This means, therefore, that while ascriptions of randomness are sensitive to the requirements of the agents who are using the concept and making the ascriptions, they are nevertheless

objectively determined, by the theories it is (objectively) appropriate for those agents to utilise. Randomness is thus an extrinsic property of events, dependent on properties of agents and the theories they use. This observation will become important below (§6.3), when discussing whether randomness as I have defined it is *subjective*.

6.1 Clarification of the Definition of Randomness

The definition of randomness might be further clarified by close examination of a particularly nice example that Tim Williamson proposed to me. Williamson's example was as follows: let us suppose that I regularly play chess against an opponent who is far superior to me. Not only does he beat me consistently; he beats me without my being aware at the time of his strategy and without my being able to anticipate any but the most obvious of his moves. I cannot predict what his moves will be. *Prima facie*, it may appear that my proposal is committed to classifying his moves as random; if true, that would pose a serious problem for the view.

Thankfully, there exist at least three lines of response to this example, each of which illuminates the thesis that randomness is unpredictability. Firstly, note that unpredictability is theory relative. It is not only the statistical aspects (i.e. actual frequencies of outcomes) of a phenomenon which dictate how it will be represented by theory; if I am convinced that my opponent is an agent who reasons and plans, no theory I will accept will have the consequence that his chess-playing behaviour is entirely random. Indeed, we will never regard these apparently probabilistic outcomes as indicative of genuine probabilistic independence (since genuine probabilities have a modal aspect not exhausted by the actual statistics). What we have in this case is not sufficient for randomness because we will never accept that the goal-directed activities of a rational agent are genuinely unpredictable, nor are those behaviours really probabilistically independent of preceding states: I certainly regard my opponent as being in a position to predict his own behaviour, and to predict it on the basis of the current state of play. Of course, in this situation, the theories which are directly available to me are not sufficient to enable me to predict that behaviour.

This leads to consideration of a second point. It is essential to note that judgements of predictability will typically be made by an epistemic or scientific community, and not a particular individual. It is communities which accept scientific

theories, and the capabilities and expertise of each member of the community contributes to its predictive powers. This is because the set of available prediction functions in a given theory does not reflect merely personal idiosyncrasies in understanding the theory, but instead reflects the intersubjective consensus on the capabilities of that theory. Since the relevant bearer of a predictive ability is an epistemic community, a phenomenon is judged random with respect to a community of predictors, not an individual. My chess playing opponent and myself are presumably members of the same scientific community and the theories we jointly accept make his chess playing predictable—he knows the theory while I accept his authority with respect to knowing it and judge his play predictable, even if not by me. This serves to reinforce the point that the ‘availability’ to me of a theory, or of a prediction function, is not a matter of what’s in my head, but rather of what theories count as normative for my judgements, given the kind of person I am and the kind of community I inhabit. One could, of course, define a concept of ‘personal unpredictability’, to capture those uses of the term ‘unpredictable’ that reflect the ignorance and incapacity of a particular individual. But—and this merely underscores the importance of the communitarian concept—such a personal unpredictability would have little or no claim to capture the central uses of the term ‘unpredictable’, nor any further useful application in the analysis of randomness or other concepts.

A third response also undermines the claim that this chess player’s moves are unpredictable. For this is exactly the kind of situation where one might be frequently surprised by the moves that are made, but one can in retrospect assimilate them to an account of the opponent’s strategy. That is, while playing I operated with a theory which was not sufficient to make accurate predictions concerning my opponent’s behaviour; in retrospect, and upon due consideration of his play, I can come to develop my understanding of that play, and hence develop better accounts of the nature of his chess playing. I can then realise that his behaviour was not random, though it may have *appeared* random at that time. Moreover, it may have been (internally) epistemically acceptable for me at the time to judge his behaviour as random (setting aside for the time being the preceding two responses), though in retrospect I can see that I had no robust external warrant for that judgement.

This last response illustrates a point that may not have been clear from the foregoing discussion: no mention was made, in the definition or its glosses, of any

temporal conditions on the appropriateness of predictive theories for agents. That is, randomness is relative to the best theory of some phenomenon, where which theory counts as best is partially dictated by the cognitive and pragmatic parameters appropriate for some community of agents. It does not, therefore, depend on whether those agents are actually in possession of that theory. Obviously it would be inappropriate to criticise past predictors on the grounds that they made internally warranted judgements of randomness that were false by the lights of theories we currently possess. On the other hand, it is true that they would deserve censure had it been the case that they were in possession of the best theory of some phenomenon, and had made judgements of predictability which were at variance with that theory. That is the sense in which theory-relative judgements of predictability are supposed to be *normative* for agents of the kind in question. As such, it is clear that contingencies of ignorance shouldn't lead us to count something as random; it is a kind of (pragmatically/cognitively/theoretically) necessary lack of predictive power which makes an event random. To turn back to the *post facto* analysis of my opponent's play: while playing I made a (perhaps) warranted judgement that it was random. But that judgement was at best preliminary and defeasible, for it is clear that it would be in principle possible for me to come to possess (or to defer to an expert's possession of) a good predictive theory of that play, and hence to recognise the sense behind what appeared wrongly to be random play. By contrast, events that are genuinely random do not contribute in this way further illumination of the process of which they are outcomes: no after the fact analysis of a random event will make greater predictive power available to me or my epistemic brethren. In one sense this is a simple corollary of the fact that the Bernoulli process is the paradigm random process, and outcomes in such a process are probabilistically, and hence predictively, independent. But in another it provides an important illustration of the application of the definition of randomness—judgements of randomness can be incorrect though warranted, and outcomes of such a process may well serve as evidence undermining the warrant for the judgement.³⁷

6.2 Randomness and Probability

One may be wondering what kind of interpretation of probability goes into the

definition.³⁸ Obviously, credences play a central role in attributions of randomness, as it is only by way of updating credences that theories yield actual predictions. As such, as long as an agent has credences that could be rationally updated in accordance with the best theory for the community of which that agent is a member—that is, whose credence function is suitably deferential to expert credence functions (van Fraassen, [1989], §§8.4–8.5)—we have the minimum necessary ingredients for potentially correct judgements of randomness. Actual judgements of randomness approach correctness as the actual updating of credences more closely approximates the updating that would be licensed by possession of the best theory. However, there is a further question concerning whether there are other kinds of objective probabilities (‘chances’) which are disclosed by the theories in question and count as normative for the credences of the predictor, via something like the Principal Principle (Lewis, [1980]). I hope that the account is neutral on this important issue, and I hope that, no matter which account (if any) turns out to be correct, it can simply be slotted into this interpretation of randomness.

In fact, the only requirement that my account of randomness makes on an interpretation of probability is that an account must be given of the content of probabilistic models in scientific theories. That is, the interpretation must explain what feature of objective probability allows it to influence credence, and to shape expectations concerning the way the world will turn out, given that all the agent does is accept some theory which features probabilistic models.³⁹ Most naturally it might be thought that an objective account of probability could meet this demand, but subjectivist accounts must also be able to do so, although perhaps less easily. Perhaps the only account that the view is not compatible with is von Mises’ original frequency view: since he includes randomness as part of the definition of probability, on pain of circularity he cannot use this definition of randomness, which already mentions probability.

Von Mises’ discussion of randomness was motivated by his desire to find firm grounds for the applicability of probability to empirical phenomena. I completely agree: random phenomena are frequently characterised by the fact that they can typically be given robust probabilistic explanations, particularly in terms of the probabilistic independence of certain events and certain initial data. But even if the grounds we have for applying probabilistic theories lie in our own cognitive

incapacities, that does not hold for the probabilities postulated by those theories. Just because predictability is partially epistemic, and hence randomness is partially epistemic, doesn't mean that the probability governing the distribution of predictions is epistemic. So our cognitive capacities and pragmatic demands lead to the suitability of treating phenomena as random, that is, modelling them probabilistically. Our epistemic account of randomness therefore provides a robust and novel explanation of the applicability of probabilistic theories even in deterministic cases, without having to mount the difficult argument that there are objective chances in deterministic worlds, and without sacrificing the objectivity of genuine probability assignments by adopting a wholesale subjectivist approach to probability. Randomness then has important metaphysical consequences for the understanding of chance, as well as being internally important to the project of understanding scientific theories that use the concept. Our epistemic stance mandates the use of probabilistic theories; the connection between the probabilities in those theories, and the credences implicit in our epistemic states, is by no means direct and straightforward.

6.3 Subjectivity and Context-Sensitivity of Randomness

I have emphasised repeatedly that predictability is in part dependent on the properties of predictors. What one epistemic agent can predict, another with different capacities and different theories may not be able to predict. Laplacean gods presumably have more powerful predictive abilities than we do; perhaps for such gods, nothing is random. Or consider a fungus, with quite limited representational capacities and hence limited predictive abilities. Almost everything is random for the fungus; it evolved merely to respond to external stimuli, rather than to predict and anticipate. It may appear, then, that judgements of predictability, and hence of randomness, must be to some extent subjective and context-sensitive. There is a worry that this subjectivity may seem counterintuitive. It may also seem quite worrying that a subjective concept may play a role in successful scientific theorising. I wish now to defuse these worries.

Firstly, it is a consequence of our remarks in §6.1 that two epistemic agents cannot reasonably differ *merely* over whether some process is unpredictable or random. If they rationally disagree over the predictability of some phenomenon, they

must be members of different epistemic communities, in virtue of adopting different theories or having different epistemic capacities or pragmatic goals. It should be quite unexceptional that agents who differ in their broader theoretical or practical orientation may differ also in their judgements of the predictability of some particular process.

Secondly, there will be reasonably straightforward empirical tests of the predictive powers of that predictor who claims the process is not random. This disagreement will then be resolved if one takes these empirical results to indicate which theory more correctly describes the world, and which therefore deserves to be adopted as the best predictive theory.

Given these qualifications, it might seem misleading to label the present account 'subjective'.⁴⁰ For, given values for the parameters of precision of observations and required accuracy of computations, and given a background theory, whether a process is predictable or not follows immediately. When we recognise that these parameter values do not vary freely and without constraint from agent to agent, but are subject to norms fixed by the communities of which agents are a part, it seems that rational agents can't easily disagree over randomness, and that purely personal and subjective features of those agents do not play a significant role in judgements of randomness. It does not seem quite right to call predictability 'subjective' simply because agents with opposed epistemic abilities and commitments may reasonably disagree over its applicability. And insofar as we remain content to classify predictability as subjective, these observations make it clear that it is a quite unusual form of subjectivity, for randomness and predictability are clearly not applied on a purely personal basis, arbitrarily and without rational constraint, and as such are capable of having further scientific significance.

But in this sense few concepts are truly subjective. Like other folk monadic properties, randomness can be analysed as a relation with the other terms fixed contextually. Consider so-called 'subjective probability', which can be analysed either as a subjective property of events, or as grounded in an objective relation between an event and an agent with certain objective behavioural dispositions. In the case of predictability and randomness, it is features and capacities of the predictor that fill in the 'missing' places of the relation. Given that, it is a matter of choice whether we decide to analyse the term as a predicate with subjective aspects, or as a relation. I

would be inclined to claim that randomness is a partially (inter-)subjective property, simply to make clear where my proposal might fall in a taxonomy of analyses of randomness, but nothing of significance really turns on this. In a similar fashion we typically use the standard subjective analysis of probability, in order to make the semantics go more smoothly and to make the connection with past uses of the term clearer.

The context sensitivity of randomness is more intriguing. Here the possibility arises that, in a given context, the relevant theoretical possibilities are delimited by a theory that in other contexts would be repudiated as inadequate. There are, I have argued (§4.3), contexts where thermodynamic reasoning is appropriate, even though that theory is false. In such contexts, therefore, a judgement of randomness may be appropriate, even though the phenomenon in question might be predictable using another theory. Perhaps when stated so baldly the context-sensitivity of randomness might seem implausible. However, randomness and predictability are only context sensitive in virtue of the fact that theory-acceptance is very plausibly context sensitive. As such, no special problem of context sensitivity arises for randomness that is not shared with other theory-dependent concepts. Furthermore, the natural alternative would be an invariantist account of randomness. Such an account would not be adequate, primarily because one would have to give a theory-independent account of randomness, and this would be manifestly inadequate to explain how the concept is used in diverse branches of scientific practice (§1).

For instance, in statistical testing we frequently wish to *design* random sequences so that they pass a selected set of statistical tests. In effect, we wish to use an effectively predictable phenomenon to produce sequences which mimic natural unpredictability by being selective about which resources (which statistical tests) we shall allow the predicting agents to use. Dembski ([1991]) sees a fundamental split here between a deterministic pseudo-randomness and genuine randomness. I reject this split: accepting it would involve a significant distortion of the conceptual continuity between randomness in deterministic theories and randomness in indeterministic theories. We certainly wish to explicitly characterise some ordered and regular outcome sequences as those a genuinely random process should avoid.

But in selectively excluding certain non-random sequences we do not thereby adopt some new notion of ‘pseudo-randomness’ that applies to the sequences that

remain. Those remaining sequences play precisely the theoretical role that random sequences play; in particular, they license exactly the same statistical inferences. Better, then, to recognise that the appropriate theory for that phenomenon, in that theoretical context, classifies those phenomena as genuinely random. Randomness by design, then, is randomness that arises from our adoption of empirically successful theories; which is to say, randomness simpliciter.

7 Evaluating the Analysis

I think the preceding section has given an intuitively appealing characterisation of randomness. The best argument for the analysis, however, will be if it is able to meet the demands we set down at the beginning of §2, and if it is capable of bearing the weight of the scientific uses to which it will be put.

I maintain that unpredictability is perfectly able to support the explanatory strategies we examined in §1. In indeterministic situations, phenomena will be unpredictable at least partly in virtue of that indeterminism. Randomness therefore shares in whatever explanatory power that indeterminism demonstrates in those cases. In deterministic cases, our account of the explanatory success of randomness must ultimately rest on something else, such as the pragmatic or epistemic features of agents who accept the probabilistic theories. Note that the proximate explanation of the explanatory success of randomness, deriving from unpredictability, remains unified across both the deterministic and indeterministic situations—a desirable feature of my proposal. Our cognitive capacities are such that, in many cases, prediction of some phenomenon can only be achieved by exceedingly subtle and devious means. As such, these phenomena are best treated as a random input to the system. The fact that these models are borne out empirically vindicates our methodology; for example, we didn't have to show that rainfall was genuinely completely ontically undetermined in order to do good science about the phenomenon in question. This is similarly the case with random mating, weather prediction, noise and error correction, and coin tossing. In random sampling (and game theory), we merely need to use processes unpredictable by our opponents or by the experimental subjects to get the full benefits of the statistical inference: if they are forced to treat

the process as random, then any skill they demonstrate in responding to that process must be due to purely intrinsic features of the trials to which they are responding.

The defining feature of the scientific theories at which we looked in §1 is the presence of exact and robust statistical information despite ignorance about the precise constitution of the individual events. Rainfall events have a definite probability distribution, but precisely when and where it rains is random. If this is the hallmark of random phenomena, then we can easily see why the particular probabilistic version of unpredictability we used to define randomness is appropriate. Indeed, in paradigm cases, unpredictability (and hence randomness) involves the probabilistic independence of some predicted event and the events which constitute the resources for the prediction. In such cases, one can easily see how the inferential strategies we have identified are legitimate. With respect to random sampling, it is the probabilistic independence of the choice of test subjects from the choice of test regimes that allows for the application of significance tests on experimentally uncovered correlations. In the case of random mating, the fact that mating partner choice is probabilistically independent of the genetic endowment of that partner that allows the standard Hardy-Weinberg law to apply. It is relaxation of this independence requirement that makes for non-random mating. This probabilistic aspect of randomness and unpredictability is crucial to understanding the typical form of random processes and their role in understanding objective probability assignments by theories.

How does our analysis of randomness as unpredictability do on our four demands (§2)?

1. **Statistical Testing** Sequences that are unpredictable to an agent can be effectively produced, since those sequences do not need to have some known genuine indeterminacy in their means of production in order to ground the statistical inferences we make using them. Subjecting the process to a finite battery of statistical tests designed to weed out sequences that are predictable by standard human subjects is, while difficult, nevertheless possible. Correlations between the test subjects and the random sequence can still occur by chance, but since there can be no a priori guarantee that could ever rule out accidental correlations even in the case of genuinely indeterministic sequences, no account of randomness

should wish to eliminate the possibility. We should only rule out predictable sources of correlation other than the one we wish to investigate.

2. **Finite Randomness** Single events, as well as finite processes, can be unpredictable.
3. **Explanation and Confirmation** A probabilistic theory which classifies some process as random is, as a whole, amenable to incremental confirmation (Howson & Urbach, [1993], ch. 14). Moreover a particular statistical hypothesis which states that the process has a unpredictable character can also be incrementally confirmed or disconfirmed, as the evidence is more or less characteristic of an unpredictable process. A special case is when the phenomenon is predicted better than chance; this would be strongly disconfirmatory of randomness. When confirmed, there seems no reason why such theories or hypotheses cannot also possess whatever features make for good explanations; they can surely form part of excellent statistical explanations for why the processes exhibit the character they do. We have gone to some lengths above to show that unpredictability can fill in quite adequately for randomness in typical uses; there seems no reason why it could not effectively substitute in explanations as well.
4. **Determinism** Unpredictability occurs for many reasons independent of indeterminism, and is compatible with determinism. Thus we can still have random sequences in deterministic situations, and as part of theories that supervene on deterministic theories. The key to explaining why randomness and indeterminism seem closely linked is that the theories themselves should not be deterministic, even if they are acceptable accounts of ontically deterministic situations.

Analysing randomness as unpredictability, I maintain, gives us the features that the intuitive concept demands, without sacrificing its scientific applicability. It certainly does better than its rivals; even without them, it captures enough of our intuitions to truly deserve the name. The final, and most demanding test would be to see how the account works in particular cases: how, for example, to cash out the

hypothesis that the mate of a female Cambrian trilobite was chosen at random from among the male members of her population.⁴¹ In outline, my proposal is that a correct account of trilobite mating would show that there is no way for us to predict (retrodict) better than chance which mate would be (was) chosen, when knowledge concerning the male individuals is restricted to their heritable properties (which of course are the significant ones in a genetic context). This entails that there is no probabilistic dependence between possession of a certain phenotype and success in mating. (Of course, given extra knowledge, such as the knowledge concerning which male actually did succeed in mating with this individual, or given facts about location or opportunity, we can predict better than chance which male would be successful; these properties are not genetic, and do not conflict with the assumption of random mating.) This account of how to apply the theory must remain a sketch, but I hope it is clear how the proposal might apply to other cases.

Acknowledgements

Thanks to audiences at Oxford, Leeds, Princeton, and the 2002 AAP in Christchurch; particular thanks to Paul Benacerraf, John Burgess, Adam Elga, Alan Hájek, Hans Halvorson, Lizzie Maughan and Bas van Fraassen for comments on earlier written versions.

Exeter College and Oxford University
Oxford
OX1 3DP
UK
antony.eagle@philosophy.oxford.ac.uk

References

Albert, D. Z. [1992]: *Quantum Mechanics and Experience*, Cambridge, MA: Harvard University Press.

- Batterman, R. & White, H. [1996]: 'Chaos and Algorithmic Complexity', *Foundations of Physics*, **26**, pp. 307–36.
- Bell, J. S. [1987a]: 'Are there Quantum Jumps?', in Bell ([1987c]), pp. 201–12.
- Bell, J. S. [1987b]: 'On the Impossible Pilot Wave', in Bell ([1987c]), pp. 159–68.
- Bell, J. S. [1987c]: *Speakable and Unsayable in Quantum Mechanics*, Cambridge: Cambridge University Press.
- Bishop, R. C. [2003]: 'On Separating Predictability and Determinism', *Erkenntnis*, **58**, pp. 169–88.
- Chaitin, G. [1975]: 'Randomness and Mathematical Proof', *Scientific American*, **232**, pp. 47–52.
- Church, A. [1940]: 'On the Concept of a Random Sequence', *Bulletin of the American Mathematical Society*, **46**, pp. 130–135.
- Dembski, W. A. [1991]: 'Randomness By Design', *Noûs*, **25**, pp. 75–106.
- Earman, J. [1986]: *A Primer on Determinism*, Dordrecht: D. Reidel.
- Fisher, R. A. [1935]: *The Design of Experiments*, London: Oliver and Boyd.
- Frigg, R. [2004]: 'In What Sense is the Kolmogorov-Sinai Entropy a Measure for Chaotic Behaviour?—Bridging the Gap Between Dynamical Systems Theory and Communication Theory', *British Journal for the Philosophy of Science*, **55**, pp. 411–34.
- Ghirardi, G. C., Rimini, A. & Weber, T. [1986]: 'Unified Dynamics for Microscopic and Macroscopic Systems', *Physical Review D*, **34**, p. 470.

Griffiths, T. L. & Tenenbaum, J. B. [2001]: 'Randomness and Coincidences: Reconciling Intuition and Probability Theory', in *Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society*.

Hartl, D. L. [2000]: *A Primer of Population Genetics*, Cumberland, MA: Sinauer, 3rd edition.

Hellman, G. [1978]: 'Randomness and Reality', in P. D. Asquith & I. Hacking (eds.), *PSA 1978*, Chicago: University of Chicago Press, volume 2, pp. 79–97.

Howson, C. [2000]: *Hume's Problem: Induction and the Justification of Belief*, Oxford: Oxford University Press.

Howson, C. & Urbach, P. [1993]: *Scientific Reasoning: the Bayesian Approach*, Chicago: Open Court, 2nd edition.

Hughes, R. I. G. [1989]: *The Structure and Interpretation of Quantum Mechanics*, Cambridge, MA: Harvard University Press.

Humphreys, P. W. [1978]: 'Is 'Physical Randomness' Just Indeterminism in Disguise?', in P. D. Asquith & I. Hacking (eds.), *PSA 1978*, Chicago: University of Chicago Press, volume 2, pp. 98–113.

Jeffrey, R. C. [2004]: *Subjective Probability (The Real Thing)*, Cambridge: Cambridge University Press.

Kolmogorov, A. N. [1963]: 'On Tables of Random Numbers', *Sankhyā*, **25**, pp. 369–376.

Kolmogorov, A. N. & Uspensky, V. A. [1988]: 'Algorithms and Randomness', *SIAM Theory of Probability and Applications*, **32**, pp. 389–412.

Kyburg, H. E., Jr. [1974]: *The Logical Foundations of Statistical Inference*, Dordrecht: D. Reidel.

Laio, F., Porporato, A., Ridolfi, L. & Rodriguez-Iturbe, I. [2001]: 'Plants in Water-Controlled Ecosystems: Active Role in Hydrological Processes and Response to Water Stress. II. Probabilistic Soil Moisture Dynamics', *Advances in Water Resources*, **24**, pp. 707–23.

Laplace, P.-S. [1951]: *Philosophical Essay on Probabilities*, New York: Dover.

Lewis, D. [1980]: 'A Subjectivist's Guide to Objective Chance', in Lewis ([1986]), pp. 83–132.

Lewis, D. [1984]: 'Putnam's Paradox', *Australasian Journal of Philosophy*, **62**, pp. 221–36.

Lewis, D. [1986]: *Philosophical Papers*, volume 2, Oxford: Oxford University Press.

Li, M. & Vitanyi, P. M. B. [1997]: *An Introduction to Kolmogorov Complexity and Its Applications*, Berlin and New York: Springer Verlag, 2nd edition.

Martin-Löf, P. [1966]: 'The Definition of a Random Sequence', *Information and Control*, **9**, pp. 602–619.

Martin-Löf, P. [1969]: 'The Literature on von Mises' Kollektivs Revisited', *Theoria*, **35**, pp. 12–37.

Martin-Löf, P. [1970]: 'On the Notion of Randomness', in A. Kino (ed.), *Intuitionism and Proof Theory*, Amsterdam: North-Holland.

Mayo, D. [1996]: *Error and the Growth of Experimental Knowledge*, Chicago: University of Chicago Press.

Montague, R. [1974]: 'Deterministic Theories', in R. H. Thomason (ed.), *Formal Philosophy*, New Haven: Yale University Press.

Popper, K. [1982]: *The Open Universe*, Totowa, NJ: Rowman and Littlefield.

Rodriguez-Iturbe, I. [2000]: 'Ecohydrology', *Water Resources Research*, **36**, pp. 3–10.

Rodriguez-Iturbe, I., Porporato, A., Ridolfi, L., Isham, V. & Cox, D. R. [1999]: 'Probabilistic Modelling of Water Balance at a Point: the Role of Climate, Soil and Vegetation', *Proceedings of the Royal Society of London*, **455**, pp. 3789–3805.

Schurz, G. [1995]: 'Kinds of Unpredictability in Deterministic Systems', in P. Weingartner & G. Schurz (eds.), *Law and Prediction in the Light of Chaos Research*, Berlin: Springer, pp. 123–41.

Shannon, C. E. & Weaver, W. [1949]: *Mathematical Theory of Communication*, Urbana, IL: University of Illinois Press.

Sklar, L. [1993]: *Physics and Chance*, Cambridge: Cambridge University Press.

Skyrms, B. [1996]: *Evolution of the Social Contract*, Cambridge: Cambridge University Press.

Smith, P. [1998]: *Explaining Chaos*, Cambridge: Cambridge University Press.

Stalnaker, R. C. [1984]: *Inquiry*, Cambridge, MA: MIT Press.

Stone, M. A. [1989]: 'Chaos, Prediction and Laplacean Determinism', *American Philosophical Quarterly*, **26**, pp. 123–31.

Strevens, M. [2003]: *Bigger than Chaos: Understanding Complexity through*

Probability, Cambridge, MA: Harvard University Press.

Suppes, P. [1984]: *Probabilistic Metaphysics*, Oxford: Blackwell.

van Fraassen, B. C. [1980]: *The Scientific Image*, Oxford: Oxford University Press.

van Fraassen, B. C. [1989]: *Laws and Symmetry*, Oxford: Oxford University Press.

van Lambalgen, M. [1987]: ‘Von Mises’ Definition of Random Sequences Revisited’, *Journal of Symbolic Logic*, **52**, pp. 725–55.

van Lambalgen, M. [1995]: ‘Randomness and Infinity’, Technical Report ML-1995-01, ILLC, University of Amsterdam,
www.illc.uva.nl/Publications/ResearchReports/ML-1995-01.text.ps.gz.

Ville, J. [1939]: *Étude Critique de la Notion Collectif*, Paris: Gauthier-Villars.

von Mises, R. [1957]: *Probability, Statistics and Truth*, New York: Dover.

Wigner, E. P. [1961]: ‘Remarks on the Mind-Body Question’, in I. J. Good (ed.), *The Scientist Speculates*, New York: Basic Books, pp. 284–302.

Endnotes

¹ Another example is more recent: ‘we say that an event is random if there is no way to predict its occurrence with certainty’ (Frigg, [2004], p. 430).

² Thanks to Steven French for emphasising the importance of these motivating remarks.

³ This is unlike the random mating example (§1.2), where we have deterministic transitions between probabilistically characterised states.

⁴ These are modelled by a Poisson distribution over times between rainfall events, and an exponential probability density function over volumes of rainfall.

⁵ Strevens ([2003]) is a wonderful survey of the way that probabilistic order can emerge out of the complexity of microscopic systems.

⁶ For details, see Albert ([1992]), Hughes ([1989]).

⁷ The law relates genotype distribution in the offspring generation to allele distribution in the parents.

⁸ See also Howson ([2000], pp. 48–51) and Mayo ([1996]). The question whether Bayesian statistics should also use randomisation is addressed by Howson & Urbach ([1993], pp. 260–74). One plausible idea is that if Bayesians have priors that rule out bizarre sources of correlation, and randomising rules out more homely sources of correlation, then the posterior after the experiment has run is reliable.

⁹ Thanks to Spencer Maughan for the example.

¹⁰ Consider some statistical test such as the χ^2 test. The probability arising out of the test is the probability that chance alone could account for the divergence between the observed results and the hypothesis; namely, the probability that the divergence between the observed sequence and the probability hypothesis (the infinite sequence) is not an indication that the classification is incorrect. A random sequence is then one that, even given an arbitrarily small probability that chance accounts for the divergence, we would not reject the hypothesis.

¹¹ A comprehensive survey of complexity and the complexity-based approach to randomness is Li & Vitanyi ([1997]). See also Kolmogorov & Uspensky ([1988]), Kolmogorov ([1963]), Batterman & White ([1996]), Chaitin ([1975]), van Lambalgen ([1995]), Earman ([1986], ch. VIII), Smith ([1998], ch. 9), Suppes ([1984], pp. 25–33).

¹² Though problems remain. The mere fact that we can give results about the robustness of complexity results (namely, that lots of universal machines will give roughly the same complexity value to any given sequence) doesn't really get around the problem that any particular machine may well be biased with respect to some particular sequence (Hellman, [1978]; Smith, [1998]).

¹³ A related approach is the so-called 'time-complexity' view of randomness, where it is not the space occupied by the programme, but rather the time it takes to compute its output given its input. Sequences are time-random just in case the time taken to compute the algorithm and output the sequence is greater than polynomial in the size of the input. Equivalently, a sequence is time-random just when all polynomial time algorithms fail to distinguish the putative random string from a real random string (equivalent because a natural way of distinguishing random from pseudo-random is by computing the function) (Dembski, [1991], p. 84).

¹⁴ For instance, Brudno's theorem establishes a connection between KCS-randomness and what is known as *Kolmogorov-Sinai entropy*, which has very recently been given an important role in detecting randomness in chaotic systems. See Frigg ([2004], esp. p. 430).

¹⁵ There are also difficulties in extending the notion to infinite sequences, but I consider these far less worrisome in application (Smith, [1998], pp. 156–7).

¹⁶ There does not exist an algorithm which on input k yields a KCS-random sequence \mathbf{S} as output such that $|\mathbf{S}| = k$, nor does there exist an algorithm which on input \mathbf{S} yields output 1 iff that sequence is KCS-random (van Lambalgen, [1995], pp. 10–1). This result is a fairly immediate corollary of the unsolvability of the halting problem for Turing machines.

¹⁷ There is some psychological research which seems to indicate that humans judge randomness of sequences by trying to assimilate them to representative outcomes of

random processes. Any product-first conception of randomness will have difficulty explaining this clearly deep-rooted intuition (Griffiths & Tenenbaum, [2001]).

¹⁸ Some complications are induced if one attempts to give this kind of account for relativistic theories without a unique time ordering, but these are inessential for our purposes (van Fraassen, [1989]).

¹⁹ There are also purported proofs of the compatibility of randomness and indeterminism (Humphreys, [1978]). I don't think that the analysis of randomness utilised in these formal proofs is adequate, so I place little importance on these constructions.

²⁰ John Burgess suggested the possibility of this response to me—and pointed out that some remarks below (particularly §§4.3 and 6.3) might seem to support it.

²¹ An infamous example of this is the bastardised notion of 'epistemological determinism', as used by Popper ([1982])—which is no form of indeterminism at all. The unfortunately named distinction between 'deterministic' and 'statistical' hypotheses, actually a distinction concerning the predictions made by theories, is another example of this persistent confusion (Howson, [2000], pp. 102–3).

²² For more on this, see Bishop ([2003]), Earman ([1986], ch. 1), Schurz ([1995]), Stone ([1989]).

²³ Note that, frequently, specification of the past macroscopic history of a system together with its present macrostate, will help to make its present *microstate* more precise. This is because the past history can indicate something about the bundle of trajectories upon which the system might be. These trajectories may not include every point compatible with the currently observed state. In what follows, we will consider the use of this historical constraint to operate to give a more precise characterisation of the current state, rather than explicitly considering it.

²⁴ We can also use shadowing theorems (Smith, [1998], pp. 58–60), and knowledge of chaotic parameter values.

²⁵ A perfect, deterministic prediction is the degenerate case where the probability distribution is concentrated on a single state (or a single cell of a partition).

²⁶ I thank Adam Elga for discussion of this point.

²⁷ As Hans Halvorson pointed out to me.

²⁸ For a start, see Jeffrey ([2004]), especially ch. 4.

²⁹ Note, in passing, that this definition does not make biased sequences any more predictable than unbiased ones, just because some outcome turns up more often. Unpredictability has to do with our expectations; and in cases of a biased coin we do expect more heads than tails, for example. We still can't tell what the next toss will be to any greater precision than the bias we might have deduced; hence it remains unpredictable.

³⁰ To connect up with our previous discussions, a *sequence* of outcomes is random just in case those outcomes are the outcomes of a random process. This is perfectly compatible with those outcomes being a very regular sequence; it is merely unlikely to be such a sequence.

³¹ At this point, it is worth addressing a putative counterexample raised by Andy Egan. A process with only one possible outcome is random on my account: there is only one event (one cell in the partition), which gets probability one, which is the same as its unconditional probability. It also counts as predictable, because all of the probability measure is concentrated on the one possible state. I am perfectly happy with accepting this as an obviously degenerate and unimportant case; recall the discussion of the trivial prediction function above (§4.4). If a fix is nevertheless

thought to be necessary, I would opt simply to require two possible outcomes for random processes; this doesn't seem ad hoc, and is explicitly included in the definition of unpredictability.

³² Compare the hierarchy of ergodic properties in statistical mechanics, where the increasing strength of the ergodic, mixing, and Bernoulli conditions serves to shorten the intervals after which each type of system yields random future events given past events (Sklar, [1993], pp. 235–40).

³³ Further evidence for this claim is provided by the fact that probabilistic independence is an all-or-nothing matter; and taking this as the definition of randomness would have the unfortunate effect of misclassifying partially unpredictable processes as not random.

³⁴ Is there ever randomness without probabilistic order? Perhaps in Earman's space invader case, it is implausible to think that any prior probability for the space invasion is reasonable—not even a zero prior. The event should be completely unexpected, and should not even be included in models of the theory. This would correspond to the event in question not even being part of the partition that the prediction function yields a distribution over. This, as it stands, would be a counterexample to my analysis, since that analysis requires a probability distribution over outcomes, and if there is no distribution, the event is trivially not random. I think we can amend the definition so as to capture this case; add a clause to the definition of predictability requiring there to be some prediction function which takes the event into consideration.

³⁵ This illuminates the common ground my proposal shares with Martin-Löf's statistical testing view of randomness. If we take the patterns to be provided by some potentially predictive theory, then failing statistical tests is equivalent to being unpredictable with respect to that theory. For the theory provides no resources to reject the hypothesis that the only structure governing the sequence is pure chance. But a potentially predictive theory will not have infinitely many concurrent predictions for a single predictor or group of predictors; so no theory can provide the resources for full Martin-Löf randomness and still remain predictive, except to creatures with computational abilities quite unlike our own. Nevertheless the spirit of the statistical test proposal remains, yet relativised to a set of statistical tests that can be actually applied to yield substantive information about the genesis and behaviour of a random process.

³⁶ Of course, if agents know their epistemic limitations, they may know of deterministic theories which can correctly account for the phenomena, but the use of which lies outside their capabilities. That is just one additional reason why randomness can correctly be assigned even in cases of perfect determinism.

³⁷ Williamson's example does point to a difficulty, however. Consider the hypothesis that our world is run by an omnipotent and completely rational deity, whose motives and reasons are quite beyond our ken, and hence our world appears quite capricious and arbitrary. If we accept such a theory, we must accept both that (i) the events in the world have reason and purpose behind them, being the outcomes of a process of rational deliberation by a reasonable agent; and (ii) that the best theory of such events that we might ever possess will classify them as random. This seems to me a genuine problem (though there is some question about its significance). One way around it might be to simply add a condition to the definition that, if the event in question is the outcome of some process of rational deliberation, it cannot be random, no matter how unpredictable it is. This proposal seems to avoid the problem only by stipulation. I

prefer, therefore, to suggest that any event which can be rationalised (as the act of a recognisably rational agent) will be predictable; and that therefore if this deity's actions are genuinely unpredictable, they are not rationalisable, and I propose cannot be seen as purposive in the way required for the example to have any force.

³⁸ Dorothy Edgington urged me to address this concern.

³⁹ Such a demand is tantamount to requiring the interpretation of probability be able to answer what van Fraassen ([1989], p. 81) calls the 'fundamental question about chance', which I take to be an uncontroversial but difficult standard to meet.

⁴⁰ John Burgess made this point.

⁴¹ I owe the question, and the illustrative example, to John Burgess.