

Best System Approaches to Chance

Wolfgang Schwarz

October 2014

Abstract. Certain physical processes such as the decay of radioactive elements seem to involve an objective, non-epistemic type of probability: *chance*. According to “best system” accounts, chance is not a fundamental physical quantity, but a kind of statistical summary of actual outcomes. The proposed interpretation explains the close connection between chance on the one hand and relative frequency, disorder, and rational credence on the other hand. It has also proved useful in understanding probabilistic theories outside fundamental physics. On the downside, the proposal arguably fails to vindicate some philosophical preconceptions about chance.

1 Introduction

In the early 20th century, developments in quantum physics suggested that the present state of the world only determines certain probabilities for future states. This type of probability is often called *chance*. Unlike degree of belief, chance does not vary from person to person. Unlike degree of evidential support, it is contingent and not relative to a body of evidence. There is an intriguing normative connection between chance and degree of belief (discussed in section 5 below), but this connection does not tell us what chance is, for we can hardly assume that the basic laws of physics involve irreducibly normative and psychological notions. So what is chance?

The question is philosophically controversial in part because chance is a modal phenomenon. The fact that an event has chance x appears to entail neither that it takes place nor that it doesn't take place; chance therefore seem to point beyond the history of actual outcomes and events to a sphere of mere possibilities. Indeed, chance looks like a graded form of physical necessity: the higher the chance of an event, the closer it is to physical necessity.

Philosophers are deeply divided on the status of modal truths in general and nomic truths in particular – truths related to physical possibility and necessity. Besides chance, prominent nomic phenomena include causation, counterfactuals, dispositions and laws of nature. One approach to the nomic, going back to Hume and further to the medieval nominalists, holds that nomic facts about what *could* or *would* or *must* are always reducible to facts about what *is*. For the most part, Humeans do not deny the reality of nomic phenomena; they agree that there are laws of nature, chances, dispositions etc. But

they maintain that these things are derivative, determined by more fundamental, non-modal elements of reality. Thus Humeans might identify laws of nature with regularities in the history of physical events, and chances with relative frequencies. This ensures that whenever two possible worlds agree in non-nomic respects, they also agree with respect to laws and chance.

On the opposite side are those who follow Aristotle and maintain that nomic elements are weaved right into the fabric of reality, as Aristotle’s “substantial forms” perhaps, or as primitive powers or primitive laws. With respect to chance, a characteristic Anti-Humean view is the primitive propensity account, defended e.g. in [Mellor 1971] and [Giere 1973] (see also Donald Gillies’ chapter in this volume). Here chance is assumed to be a basic physical quantity not unlike mass or charge. Just as an atom’s mass grounds a disposition to interact in certain ways with its environment, an atom’s propensity to decay within a certain interval of time grounds a “partial disposition” to decay within that time.

Humeans and Anti-Humeans are generally easy to identify, although it is hard to state their disagreement precisely, without relying on controversial physical and metaphysical assumptions. Humeanism is sometimes characterized (e.g. in [Lewis 1986b: ix–xi], [Lewis 1994: 225f.] and [Loewer 2012: 116]) as the view that all truths supervene on the distribution of fundamental categorical properties over points or regions of spacetime (or some other physically basic space), where a categorical property is one whose instantiation in a region entails nothing about the instantiation of fundamental properties outside that region. However, this would not exclude primitive propensities, since instantiation of a propensity does not strictly entail anything about other times or places.

The present article is concerned not with Humeanism in general, but with a particular Humean analysis of chance: the “best system account”, first proposed by David Lewis in [Lewis 1986b: 128] and [Lewis 1994]. The analysis can be seen as a sophisticated descendant of the frequency interpretation (see La Caze’s article in this volume). Unlike simple frequentism, the best system account allows chances to come apart from actual frequencies: if there is a 50 percent chance that radium atoms decay within 1600 years, it does not follow that exactly 50 percent of the atoms really decay within that time (and therefore that there is an even number of radium atoms). Rather, chances are identified with probabilities in ideal physical theories whose aim is to provide a kind of summary statistic of actual outcomes; getting close to the frequencies is one virtue of probabilistic theories, but it trades off against other virtues such as comprehensiveness and simplicity.

2 Laws and chance

Lewis’s best system account of chance begins as a best system account of laws, also known as the Mill-Ramsey-Lewis account, due to its origins in [Mill 1843], [Ramsey 1978] and [Lewis 1973]. If we look at popular theories in fundamental physics – the theory

of general relativity, say, or the standard model of particle physics – we find that they predict a staggering variety of facts with very high precision, based on relatively few and simple assumptions. That is no coincidence. What physicists aim for is precisely to find simple principles that allow a unified explanation for all sorts of complex phenomena. Of course our present theories may turn out to be false; some physicists also hope to find an even more unified and comprehensive “theory of everything”. Suppose there is such a theory – whether or not we will ever find it. Suppose it makes no false predictions, and there is no other theory that accounts for an even greater range of facts with even simpler rules. On the best system account, it follows that the rules of this best theory are the true laws of nature. Lewis [1994: 231f.] succinctly states the general analysis:

Take all deductive systems whose theorems are true. Some are simpler, better systematized than others. Some are stronger, more informative, than others. These virtues compete: an uninformative system can be very simple, an unsystematized compendium of miscellaneous information can be very informative. The best system is the one that strikes as good a balance as truth will allow between simplicity and strength. [...] A regularity is a law iff it is a theorem of the best system.

In worlds like ours, there are especially salient regularities linking earlier and later states of physical systems. Perhaps these regularities can be captured by simple equations expressing a functional relationship between the state of isolated systems at one time and their past and future evolution. But perhaps things are not so simple. Perhaps a physical system in state X sometimes evolves into state Y and sometimes into Z , without any intrinsic difference in the initial states or their histories. In this case, it might nevertheless be useful to learn that X states are followed by either Y states or Z states. Moreover, it might be valuable to know something about the proportions: is one of the outcomes much more common than the other? A good theory might then specify probabilities for a system in state X to turn into Y or Z . In general, if the history of a world reveals a pervasive but “noisy” dependence between two quantities, then a good theory might specify a probabilistic connection between these quantities to convey information about their distribution. Lewis suggests that these probabilities are the chances: “the chances are what the probabilistic laws of the best system say they are” [1994: 234].

This may look circular, but it is not. The idea is that one can evaluate probabilistic theories without first interpreting their probabilistic statements:

As before, some systems will be simpler than others. Almost as before, some will be stronger than others: some will say either what will happen or what the chances will be when situations of a certain kind arise, whereas others will fall silent both about the outcomes and about the chances. And further,

some will fit the actual course of history better than others. That is, the chance of that course of history will be higher according to some systems than according to others. [...] The virtues of simplicity, strength, and fit trade off. The best system is the system that gets the best balance of all three. [1994: 234]

To illustrate, imagine a world consisting of nothing but a short sequence of binary events (“coin tosses”), with outcomes 0111000011. Let T_1 be a theory according to which the events are independent with both 1 and 0 having constant probability $1/2$. Theory T_2 assigns probability $9/10$ to 1 and $1/10$ to 0. The total sequence then has probability $(1/2)^{10} \approx 0.001$ according to T_1 and probability $(9/10)^5(1/10)^5 \approx 0.000005$ according to T_2 . In general, the closer a theory’s probability assignments match the relative frequencies, the greater the probability of the actual sequence and therefore the better the theory’s fit. Another theory, T_3 , might relax the independence assumption and say that the probability of 1 is $2/3$ immediately after an occurrence of 1 and otherwise $1/3$. The sequence then has probability around 0.002; but this increase in fit comes at a cost in simplicity.

The real value of probabilistic theories only becomes apparent in more complex worlds. Imagine we knew the precise decay times for every atom in the history (past and future) of the world. How could we summarise this information? Listing the actual frequencies for every type of atom and every interval of time would be as unwieldy as listing the individual decay times. One could instead give the average decay time for each type of atom, or the average and the standard deviation. Much better, one could adopt the language of probability and specify an exponential distribution of decay times for each element, with the understanding that the actual frequencies approximately fit that curve. This would not allow reconstructing the exact actual pattern of decay times, because many alternative patterns would be summarised by the very same distribution, but it would convey a lot of information about the pattern in a very brief and elegant manner.

To get a full analysis of chance, we would now have to spell out the ordering that determines which theory counts as best relative to a possible history of events. It is not crucial to the best system approach how exactly these details are filled in. Nevertheless, it is worth having a closer look at some issues this involves.

3 Comparing theories

The best system account defines the chances at a world as the probabilities that figure in the best theory for that world. To this end, Lewis identifies theories with arbitrary deductively closed sets of sentences (“systems”) in some fixed language L . To avoid the choice of a language, it is tempting to instead identify theories with sets of models. But this might cause difficulties with probabilistic theories: what should count as a model

of a probabilistic theory, given that the probabilities are as yet uninterpreted? I will therefore stick with Lewis's somewhat anachronistic syntactical conception of theories.

Lewis mentions three criteria of goodness for such theories: simplicity, strength and fit. Let's take these in turn, beginning with simplicity. If theories are sets of sentences, simplicity is naturally understood as syntactical complexity. Roughly speaking, the fewer axioms are needed to generate a theory, and the simpler their logical and mathematical form, the simpler the theory.

Evidently, any such measure is highly sensitive to the chosen language. If T is a complicated (non-probabilistic) theory axiomatised by 147 long and unrelated principles, we can define an atomic predicate F to be true of an object x iff x exists in a world where those 147 principles are true; T can then be translated into the very simple theory 'everything is F ' (see [Lewis 1983: 42]). Hence Lewis requires that before entering the competition, all theories must be translated into a common language L in which all basic non-logical terms stand for fundamental properties (or magnitudes) like mass or charge. Gerrymandered predicates like F are forbidden.

This move has raised some eyebrows (see e.g. [van Fraassen 1989: 52–55], [Loewer 1997: sec.3], [Cohen and Callender 2009]). For how do we know which properties are fundamental? If two theories postulate different sets of fundamental properties, how can we tell which is right? Lewis holds that fundamental properties are also related to similarity and duplication: two objects are perfect duplicates iff they agree in the distribution of fundamental properties and relations over their parts (see [Lewis 1986a: 61f.]). But arguably this does not rule out the skeptical possibility that scientists might settle on apparently simple regularities that are not the true laws because they are not very simple any more when translated into the language of objectively fundamental properties.

Uncomfortable with this consequence of Lewis's proposal, several advocates of the best systems account have suggested an alternative on which every theory may provide its own inventory of fundamental properties (see [Loewer 2007b], [Cohen and Callender 2009]). On this view, there is no objective, theory-independent standard for comparing systems with different inventories; hence there is nothing objectively wrong with a system that encodes many (or even all) truths in a single statement 'everything is F '. Perhaps the problem with this system is only that it does not strike us as very perspicuous.

What shows up in this discussion is the other great divide in the metaphysics of science, between scientific realism and anti-realism. Lewis assumes that the world has an objective, mind-independent structure, given by the distribution of fundamental properties and relations. Scientific theories try to uncover patterns in this structure, but there is always a skeptical possibility that they fail, that the lines drawn by science do not carve nature at its joints. Cohen and Callender, on the other hand, reject the idea of objectively natural or fundamental properties: objectively, the negatively charged particles have no

more in common than an electron, a goat, and the Atlas mountains. Our preference for theories expressed in terms of charge, mass, etc. therefore only reflects idiosyncratic features of our interests and upbringing. The best system account can be made to fit both the realist and the anti-realist picture.

Lewis's second criterion is strength. An obvious way to understand strength is in terms of excluded possibilities; the more possibilities a system excludes, the greater its strength. It is not clear what exactly this means given that the excluded possibilities are typically infinite. There are also problems with probabilistic theories, which don't seem to exclude anything – especially if the probabilities are uninterpreted. On reflection, the whole idea is on the wrong track anyway. Consider two worlds w_1 and w_2 ; w_1 contains a lot of F s, all of which are G ; w_2 is like w_1 except that it contains very few F s (perhaps none). In this case, 'all F s are G ' is a better candidate to be a law of w_1 than a law of w_2 , since it provides much more valuable information about w_1 . It follows that strength is not an intrinsic aspect of theories. Whether one theory is stronger than another depends on which world is under consideration. Roughly speaking, the more F s there are in a world, and the fewer G s, the stronger is a law 'all F s are G s' (or 'the probability of an F being G is x ') with respect to that world.

John Earman has often pointed out (e.g. in [Earman 1984], [Earman 1986], and [Earman and Roberts 2005]) that the strength scientists value in physical theories is actually more specific. Modern physical theories typically provide differential equations which, combined with suitable boundary conditions concerning (say) a system's state at a particular time, allow computing the system's state at earlier and later times. What we value is strength in these dynamical laws. The ideal limit are deterministic laws, which allow complete reconstruction of a system's history from information about a single time. By contrast, strength with respect to boundary conditions is much less of a virtue, even if it would make a system more informative. For an extreme example, imagine a world with just a few thousand particles and deterministic dynamical laws. For each time, there is a statement that describes the exact state of all particles at that time. Adding the simplest among those statements to the dynamical laws yields a maximally informative system. Nevertheless, we are not inclined to deem this a very good physical theory, nor would we say that in such a world, all facts are physically necessary.

Lewis explicitly mentions that a theory may involve statements of particular fact, in which case only the "regularities" it entails count as laws (see [Lewis 1986b: 123]). He does not explain what makes something a regularity (obviously this is not a matter of syntax), but we can assume that boundary conditions are excluded. To some extent, it is then unimportant for the analysis of laws and chance if systems may include boundary conditions or if instead their strength is measured by their informativeness when combined with external boundary conditions. Either way, the boundary conditions do not come out as laws or as physically necessary.

One may still wonder why a good system should include dynamical laws at all, or dynamical laws plus boundary conditions. If all that counts is informativeness or logical strength that is hard to explain. But we've seen that informativeness by itself is not the right standard to begin with. The right standard is world-relative, and – well – it privileges strength in dynamical laws. On the best system account, the standards *define* the notions of laws and chance. To the question ‘what makes those standards right?’ there is no substantive, metaphysical answer. We could have used a different notion of (quasi-)laws defined by other standards. That in fact we are more interested in systems with strong dynamical laws is easy to understand: we often have information about the present (or recent past) of physical systems, and want to know more about their future and distant past. A system of laws that applies (approximately) to (more or less) isolated subsystems of the universe, in all sorts of boundary conditions, is a very useful thing to know.

Lewis's third criterion for theories is “fit”. Lewis suggests that this is measured by the probability a theory assigns to the entire history of the world. This is a natural idea, but it suffers from a technical difficulty, known as the *zero-fit problem*: good theories often assign zero probability to the entire history of a world, either because the world contains infinitely many chance events, or because the space of outcomes for individual events is infinite. For example, if the decay probabilities for radium atoms follow an exponential distribution, the probability of any particular atom decaying at any particular time is zero. Lewis [1980] did regard this as a merely technical problem, pointing at [Bernstein and Wattenberg 1969] for a solution based on probability theories that allow for infinitesimal values. Elga [2004] argues that these infinitesimals cannot be relied upon to yield the desired ordering on theories in terms of fit. Drawing on [Gaifman and Snir 1982], he instead suggests to define fit by looking at the probabilities a theory assigns to a carefully chosen set of truths about a history, rather than the conjunction of all truths.

Another problem with Lewis's (and to some extent Elga's) proposal is that a good theory may assign no probability at all to the total history of the world. A stochastic differential equation, for example, directly specifies probabilities only to transitions between states. If the world has an initial state, one could perhaps compute a probability for the entire history relative to this starting point; but this is hardly a general solution.

At this point, it might be a good idea to look at “goodness of fit” measures developed in statistics, such as the chi-squared test of classical statistics or one of its Bayesian counterparts. Interestingly, these measures usually include simplicity constraints to avoid overfitting. This illustrates that Lewis's three criteria need not be treated as independent: the ordering that defines the best system doesn't need to come about by somehow balancing three independent scores. In the context of the best system approach, an especially noteworthy approach to statistical model selection is the framework of “Minimum Description Length”, introduced in [Rissanen 1978]. Lewis's Humean perspective on the

relationship between theories and the world bears a striking resemblance to Rissanen’s perspective on the relationship between models and data. Rissanen rejects the idea that the goal of statistical inference is to find the true probability distribution that “generated” the data (see especially [Rissanen 1989]); likewise, Lewis rejects the idea that the pattern of events in the world is “generated” by some hidden, underlying probability distribution; objective probabilities simply state noisy regularities in the pattern of events.

Scientific practice also suggests further conditions on good theories. For example, physicists generally prefer theories with fewer primitive constants, even if reducing the constants comes at a slight cost in simplicity. To measure this, it can be useful to consider not only an individual theory, but the parametric family of theories that results by leaving the value of constants open and checking how many patterns of outcomes can be made to fit by tweaking the constants: the more, the worse (see [Myung et al. 2000]).

However, we should not assume that the standards for good theories relevant to the best system account precisely match the standards used in science (although Lewis says so in [1986b: 123]). Model selection criteria are often constrained by computational complexity. Scientists also tend to prefer theories that are conservative with respect to earlier theories, presumably on the ground that it is much easier to see things correctly if one is standing on the shoulders of giants. These features seem out of place in the analysis of laws and chance. When scientists invoke criteria of statistical model selection, they consider the adequacy of theories in the light of a limited set of observations, with an eye to predicting the unobserved. By contrast, for the best system account we compare a theory’s probabilities with the total history of the world (observed and unobserved), in order to define a goodness order on the space of theories which then defines the notions of laws and chance. The only direct constraint here is the plausibility of the resulting analysis. To which we turn now.

4 Playing the chance role

Best system accounts identify chance with certain features of the total history of events in a world. It is relatively uncontroversial that these features exist. But can they play the role of chance in scientific and philosophical thought?

The question, it turns out, is ill-posed. There is no such thing as *the* role of chance in scientific and philosophical thought. A symptom of this is the debate over whether chance is compatible with determinism: if the laws of nature are deterministic, can there still be non-trivial chances pertaining for instance to the outcomes of coin tosses? Some, including Popper [1982: 105], Lewis [1986b: 120], Hájek [1997: 221] and Schaffer [2007], take the answer to be obviously ‘no’. Others, including Levi [1990], Loewer [2001] and Hofer [2007] deem it equally non-negotiable that the answer is ‘yes’. Philosophers also disagree over whether chance is essentially dynamical, linking earlier to later states, and

whether past events can still be chancy. One can hardly expect any one analysis of chance to satisfy all these contradictory constraints.

Some aspects of chance, at least, are uncontroversial. First of all, most authors agree that chance (in the sense of physical probability) satisfies some form of the probability axioms. On the best system account, this could simply be the consequence of a stipulation that good systems should not only be logically consistent, but also probabilistically coherent. More ambitiously, one might argue that it follows from the fact that for every probabilistically incoherent theory there is a coherent theory that is guaranteed to have greater fit no matter what the world is like (see [Joyce 2009] for observations of this type, although not applied to chance).

Second, it is usually agreed that chances figure in laws of nature and are thus linked to other nomic phenomena such as causation and counterfactuals. By itself, the best system account of chance does not establish this; however, the account naturally goes with a corresponding account of laws and other nomic phenomena, which secures the desired links.

Scientific practice requires a close connection between chance on the one hand and symmetries and frequencies on the other. In particular, relative frequencies in long series of trials are assumed to be close to the chances. This substantive fact is sometimes mistaken as a consequence of the Laws of Large Numbers; but the Laws hold for all probability functions whatsoever, even ones that are completely out of tune with the frequencies. On the best system account, the alignment between chance and frequencies is not too surprising, since a good system must have high fit. Similarly, invariance under symmetries is a plausible virtue of physical theories, so it is unsurprising that the chances should respect it.

This third feature of chance leads to a fourth: chances can be discovered by the methods of science. Again, the best system account makes this understandable, since one can reasonably hope that the methods of science get us close to the ideal theories whose probabilities are identified with the chances. Of course, there is no logical guarantee of success. Demanding as much would be implausible anyway. It is logically possible that our observations are systematically skewed, that relative frequencies in observed samples are radically different from the overall frequencies, and so on. If these skeptical possibilities show that science cannot discover the chances (on the best system account), they show that science cannot discover general truths at all.

Two further important aspects of chance concern the distribution of chance events in the world. First (or rather, fifth), the chance of an outcome is typically constant across intrinsically similar situations: if two systems perfectly agree in the distribution of mass, charge, spin, etc., we expect them to also agree about the dynamical chances (see [Arntzenius and Hall 2003]). On the best system account, this makes sense: if the chances varied widely and unsystematically between intrinsically identical situations, they

could hardly be specified in a simple and elegant theory. On the other hand (sixth), we expect disorder in the *outcomes* of chance: chance events usually have a random-looking distribution. This is also predicted by the best system account: if the outcomes of a process come in a conspicuous pattern (such as 1010101010...), a good theory would not need to resort to probabilities.

All these issues await more detailed investigation, spelling out under what conditions and on what background assumptions the best system account can vindicate the stated facts about chance. But it is fair to say that the prospects look considerably better than e.g. on the primitive propensity account. If chance is a primitive quantity logically independent of the distribution of other quantities, then it is rather mysterious – to put it mildly – why this quantity should be linked to frequencies and symmetries among other quantities, why it is doesn't vary when those other quantities are held constant, or why chance events come in disorderly patterns.

A further advantage of the best system approach is that it can vindicate many of the more controversial views about chance – even when these appear to contradict one another. This is because it allows for different species of chance. Let's reserve the title *fundamental dynamical chance* for probability functions in a Lewisian best system that pertain to later states of a world given an earlier state. Fundamental dynamical chance only exists in indeterministic worlds. Since it is always conditional on a given (earlier) state of the world, it appears to evolve over time by conditioning on history, and can be stated in the form of “history-to-chance conditionals” – just as for example Lewis [1980] and Schaffer [2007] intuit. But the best system approach also allows for other ways in which probabilities could find themselves in a best system. For example, a best system might involve a non-dynamical probability distribution over initial states of the universe, as in some versions of statistical mechanics or Bohmian mechanics. This type of chance is compatible with deterministic dynamical laws. Whether or not it deserves the honorific label ‘chance’, it is an objective physical quantity, just as Loewer [2001], Hoefer [2007] and others insist.

The best system account can even vindicate many claims of the propensity interpretation. As we saw, in well-behaved worlds with fundamental dynamical chances (in the sense of the best system account), the chances are plausibly *intrinsic to experimental setups* insofar as duplicating an experiment also duplicates the chances. Since the chances are given by “partial laws” which, instead of saying that all F s are G merely assign a probability to G given F , one might say that the condition F *partially necessitates* manifestation of G , so that the relevant system has a *partial disposition* to produce G . Like propensities, fundamental dynamical chances are *forward-directed* and *essentially conditional*. The only point where the best system account must depart from propensity accounts is when propensities are declared metaphysically fundamental.

Three more aspects of the chance role deserve special attention. One is the link between

chance and rational belief; another is the apparent independence of chances from actual outcomes. These will be discussed in section 5 and 6, respectively. The third is the objectivity of chance.

This is often regarded as a problem for best system accounts. Chances (and laws of nature) are supposed to be objective, mind-independent features of the world. But don't the standards for good theories, in particular the standards of simplicity, depend on us? Humeans with anti-realist or pragmatist sympathies might endorse this consequence and reject the idea of mind-independent laws and chances as targets of scientific inquiry. As a philosopher of a more realist bent, Lewis [1994: 232f.] took the objection more seriously. He offered two lines of response. First, he argued that it is not just a matter of taste which theories we deem more or less simple (when translated into the language of fundamental properties). Indeed, we think simpler theories are more likely to be true; this would make no sense on the assumption that the relevant standards of simplicity merely reflect epistemically irrelevant facts about human psychology. In addition, Lewis suggested that our concepts of laws and chance only have clear application in worlds where a unique best system is significantly ahead of the competition, so that the laws and chances do not depend on fine details or the comparison.

But this does not fully resolve the worry. Even if there are objective standards for what to believe about the world given limited observations, it does not follow that there are objectively privileged standards for how to best summarise the entire history of the world. Consider again our preference for strength in dynamical laws. This arguably reflects our epistemic limitations and interests. Does this not undermine the objectivity of laws and chance? The worry is easy to sense, but harder to make precise. After all, the best system account does not define chance indexically as the probability in whatever theory comes out best according to our standards, whatever they might be. Rather, a fixed set of standards is used in the definition. To be sure, these standards come from us – but where else should they come from? The same is true for all terms. Whether something falls in the extension of 'positively charged' (say) also depends on contingent facts about how we apply those words; it does not follow that positive charge is not an objective, mind-independent matter. However, there is a difference. On the realist picture, our notion of positive charge picks out a metaphysically privileged, fundamental joint in nature. Not so our notion of chance, if the best system account is right: it picks out something objective, but nothing *objectively special*. Chance is special only to creatures like us.

5 Chance and credence

Perhaps the most striking feature of chance is its connection to rational belief. As a first stab, the connection might be expressed as follows. If Cr is a rational credence function

and $Ch_t(A)=x$ a proposition specifying that the chance of A at time t is x , then

$$Cr(A/Ch_t(A)=x) = x,$$

provided $Cr(Ch_t(A)=x) > 0$. For example, knowing that a given coin has chance 1/2 of landing heads, one should have equal confidence in heads and tails. Many variations of this principle have been put forward, under various names. A problem with the present formulation is that it doesn't take into account the possibility that the agent has further information about the relevant proposition A , as when she has already seen the coin land heads (and t is a time in the past). On the other hand, if the coin toss lies in the future, it is hard to imagine any evidence that would affect one's rational confidence in heads and tails, given knowledge that the coin is fair. Taking these two observations into account leads to a version of the chance-credence link which Lewis [1980] dubbed the *Principal Principle*. Again, there are slightly different ways of stating the principle; Lewis himself gave several non-equivalent formulations. The first and best known goes in essence as follows. If Cr is any rational *initial* credence function, $Ch_t(A)=x$ a proposition saying that the chance of A at t is x , and E arbitrary *admissible* information, then

$$Cr(A/Ch_t(A)=x \wedge E) = x,$$

provided $Cr(Ch_t(A)=x \wedge E) > 0$. "Admissible" information is the kind of information one can reasonably acquire ahead of the relevant time t .

The Principal Principle, in some form or another, is crucial to our understanding of chance. It also serves as a powerful touchstone for candidate interpretations: chance can be identified with an objective quantity X only if it is plausible that X guides rational credence in the way demanded by the Principle. This is where Lewis, until the early 1990s, saw a fatal problem with his best system account – for he had discovered that the analysis is incompatible with the Principal Principle (see [Lewis 1980: 111f.] and [Lewis 1986b: 129–131]). He famously called this the "Big Bad Bug" for his whole Humean metaphysics (see [Lewis 1986b: xiv]).

The problem can be illustrated in a simple coin toss universe. Imagine a world consisting of 10 binary events, and consider again the theory T_1 that takes the outcomes to be independent with constant probability 1/2. The sequence 1111111111 then has probability $(1/2)^{10}$, just like 0111000011. But 1111111111 is not best systematised by T_1 . So T_1 assigns positive probability to situations that are incompatible with T_1 being the best system. In general, no theory on which the chance of all-1 is less than 1 best systematises the all-1 history. Thus for any x strictly between 0 and 1, $Cr(\text{all-1}/Ch(\text{all-1})=x) = 0 \neq x$.

The phenomenon at issue is best known from so-called "undermining futures". Let Z be a proposition according to which most radium atoms from now on decay within a few years. Suppose this proposition, although false, has positive chance x . In worlds where Z is true, the history of actual events may well be different enough from the actual history

to give rise to different chances (on the best system account). In particular, the chance of Z in such worlds may not be x . On the supposition that the chance of Z is x it then follows that Z is false. But the Principal Principle requires that Z has positive credence x on this supposition.

Undermining problems arise because for Humeans, information about chance contains statistical information about the history of outcomes, and everyone agrees that such information is not in general admissible. To incorporate this fact, Lewis [1994] and Hall [1994] (independently) suggested a revised version of the Principal Principle, which I state in a form due to [Hall 2004]: if Cr is a rational initial credence function, $Ch_t=f$ the proposition that the chance function at t is f , and E arbitrary information, then

$$Cr(A/Ch_t=f \wedge E) = f(A/Ch_t=f \wedge E),$$

provided $Cr(Ch_t=f \wedge E) > 0$. Roughly speaking, the idea is to treat chance as an expert, so that conditional on the information that $Ch_t=f \wedge E$ one should align one's beliefs with the expert's opinion *conditioned on that same information*. Hall [1994, 2004] argues that the new formulation is preferable to the old one on independent grounds, even if (like Hall) one rejects Humeanism about chance.

The revised Principle avoids the undermining problems. In the case of the undermining future Z , $Ch_t=f$ is logically incompatible with Z given $f(Z) < 1$, in which case $Cr(Z/Ch_t=f) = f(Z/Ch_t=f) = 0$, as desired. More illuminating is perhaps the coin toss example, because it is easier to work through in detail. Let f_1 be the chance function according to T_1 , and consider the binary sequences of length 10 compatible with the hypothesis that T_1 is the best system. Plausibly, these sequences all have five 0s and five 1s in a more or less random-looking distribution. Let H_1 be the set of these sequences. If the credence function Cr assigns equal probability to the members of H_1 , it follows that $Cr(\omega/Ch_t=f_1) = f_1(\omega/Ch_t=f_1)$ for every sequence ω . This is because $Ch_t=f_1$ rules out all sequences outside H_1 , and both f_1 and Cr are uniform within H_1 .

A potential downside of the revised Principle, apart from its unintuitive complexity, is that it requires chance to be defined over a very rich set of propositions, including propositions of the form $Ch_t=f \wedge E$. Some advocates of the best system account have therefore preferred to stick with a weakened form of the old Principle. On this view, the old Principle usually holds to a good approximation and occasionally breaks down for artificial propositions like Z or in artificially simple scenarios like that of the coin tosses (see [Hoefer 2007], [Schwarz 2014]).

Having defused the Big Bad Bug, Humeans are wont to turn the tables. Granted, Humean accounts are in tension with simple formulations of the chance-credence link, but non-Humean accounts seem to render any such link entirely unintelligible (see e.g. [van Fraassen 1989: 80–86], [Lewis 1986b: xvf.], [Lewis 1994: 239], [Loewer 2004: 1120]). Suppose, following advocates of primitive propensities, that there is a basic physical

quantity besides spin, mass, charge, etc. which maps propositions (or pairs of propositions) to numbers. Across logical space, the values of this quantity are independent of other facts about the world: in some worlds, high values generally go with true propositions, in others with false ones, and so on. Such a quantity might well exist; indeed, there might be many of them, assigning different numbers to propositions, all independent of their truth. But then how could it be a basic constraint of rationality to believe propositions to the extent that a particular one of these quantities assigns them high numbers? Why should credence follow this quantity rather than one of the others? Why take high numbers as indications of truth rather than low ones? Note that we have absolutely no evidence that our world is one where high numbers tend to go with true propositions, for we have no direct way to inspect this primitive quantity; physicists estimate chances by *using* the Principal Principle, inferring (for example) that frequently occurring outcomes must have high chance.

By contrast, if chance is a Humean quantity reflecting patterns in the distribution of actual outcomes, it is not at all mysterious that information about chance should constrain rational belief about outcomes. To be sure, it is not immediately obvious that this constraint takes the specific form of the Principal Principle (suitably weakened or revised to account for undermining), and some Anti-Humeans have expressed skepticism that Humean accounts can really explain the Principle (see [Black 1998], [Strevens 1999], [Hall 2004]). But we have already seen a simple illustration of how this might go. In the coin toss example above, we have *derived* (central instances of) Hall's revised Principal Principle from the assumption that equal initial credence goes to random-looking sequences with a fixed ratio of 0s and 1s. This is a plausible constraint on rational credence, independent of the interpretation of chance. Observations of this kind are generalised in [Schwarz 2014] (see also [Loewer 2004: 1122f.], [Hofer 2007: sec.5] for related arguments).

6 Anti-Humean intuitions

For Humeans, the history of occurrent events in a world is metaphysically prior to laws and chances. Anti-Humeans, on the other hand, often hold that laws and chances “produce”, “guide” or “govern” the history of events and therefore must have metaphysical priority. As an objection to Humeanism, this line of thought has proved rather elusive, as it is not clear what the metaphors of production, guidance or governance are supposed to express.

Anti-Humeans sometimes try to cash out the objection by pointing at links between laws and chance on the one hand and explanation, counterfactuals or causation on the other. From an Anti-Humean perspective, possible worlds without primitive nomic features are lawless worlds in which any regularity in the history of events is but a

mere coincidence; hence there can be no explanation of why (in such a world) a glass shatters when hit by a rock, and nothing definite can be said about what would have happened if, counterfactually, a glass had been hit by a rock. This is an alien and unappealing picture. But of course it is not at all the Humean picture. For example, on Lewis's Humean account of counterfactuals, causation and explanation (see [Lewis 1986b]), Humean laws do support counterfactuals, are involved in causation, and figure in causal explanations. What is revealed by the present objection is merely that Humeanism is a package deal: Humeans about laws and chance had better not accept Anti-Humean accounts of counterfactuals, causation, explanation, etc.

Arguably the best way to motivate the Anti-Humean position is to emphasise the intuitive gap between the *is* of occurrent events and the *ought* of laws and chance. Lange [2000: 48–51] describes a possible world in which there is a single particle, traveling at constant velocity. Intuitively, such a world could be governed by Newton's laws, but also by various alternatives; the actual events are too sparse to settle the question. This kind of thought experiment is especially powerful in the case of chance (see e.g. [Tooley 1987: 42–47]). Imagine a world with nothing but a thousand “coin toss” events: 517 heads and 483 tails. Does it follow that the chance of heads on each toss is exactly 0.517? Intuitively not; the sequence is quite compatible with the chance being 0.5 or 0.6 or even 0.9. But then the chances are not determined by the actual outcomes.

To suggest otherwise seems to conflate improbability with impossibility. If the chance of heads is 0.9, getting 517 heads in 1000 tosses is improbable; but improbable things can happen. In general, the hypothesis that a proposition *A* has positive chance should always be compatible with *A* being true. This is a version of what Bigelow et al. [1993] call the *Basic Chance Principle*. It is also a special case of Lewis's original Principal Principle: in the absence of inadmissible evidence, one should assign positive credence to *A* given that *A* has positive chance. On the best system account, by contrast, world histories with comparatively low chance can be ruled out *a priori*!

Humeans might try to escape these objections by denying the move from conceivability to metaphysical possibility, appealing to [Kripke 1980]. This has not been a popular response. Friends of the best system account generally take the account to provide an analysis or explication of the concept of chance, rather than an empirical hypothesis about its essence. In fact, Humeans are generally skeptical about essences and other substantive kinds of necessity and possibility. The usual reaction is therefore to bite the bullet and concede that the best system account is incompatible with some intuitions about chance. This is sometimes combined with an allegation that the intuitions come from dubious “theological” sources and hence should not carry much weight in the first place (see e.g. [Loewer 1997], [Beebe 2000]). One might also argue that the intuitions arise from uncritical applications of the Principal Principle. But one need not be so dismissive. One can accept that the Anti-Humean intuitions are part of our ordinary

conception (or conceptions) of chance. Humean accounts cannot satisfy this aspect of the chance role. But sometimes perfection is impossible. Humean accounts at least satisfy many other aspects of the chance role, including those that actually matter to science – and they do so precisely because they deny the independence of chances from actual outcomes.

7 Probability in science

When philosophers talk about chance, they mostly have in mind fundamental dynamical chance in the sense of section 4: a forward-looking probabilistic quantity that figures in dynamical laws of fundamental physics. It is doubtful whether this kind of chance exists in worlds like ours. The dynamical law of standard quantum physics, the Schrödinger equation, is deterministic. In general, the formalism of standard quantum mechanics does not seem to involve probabilities at all. (Wave function intensities are sometimes called probabilities, but interference effects preclude any direct interpretation as probabilities of underlying physical states.) Probabilities do figure in some rivals to standard quantum mechanics, notably in de Broglie-Bohm mechanics and the “GRW” theory of [Ghirardi et al. 1986] and [Pearle 1989]. The probabilities in GRW even look like fundamental dynamical chances, but those in de Broglie-Bohm mechanics do not, since they are not dynamical.

This highlights an advantage of the best system account that we have briefly touched upon in section 4: unlike e.g. the primitive propensity account, the best system account can accommodate objective physical probabilities that do not take the form of fundamental dynamical chances. Many advocates of the best system approach argue that it can serve as a general analysis of objective empirical probabilities, including for example the probabilities in de Broglie-Bohm mechanics.

Outside fundamental physics, probabilities are found all across science – from statistical mechanics to theories of genetic mutation and evolution. Here, too, the probabilities rarely look like propensities. In statistical mechanics, for example, they primarily pertain to different microphysical realisations of thermodynamic states; in coalescence theory, they are even “backward-looking”, pertaining to the time in the past at which genes in the present generation had their latest common ancestor. On the other hand, these probabilities satisfy many of the conditions discussed in sections 4 and 5, such as the connection to frequencies, randomness and rational credence. This suggests that they might be interpreted along the lines of the best system account.

The most ambitious and detailed development of this idea is the Albert-Loewer account of statistical mechanics (see [Albert 2000], [Loewer 2007a], [Loewer 2012]). Albert and Loewer argue that the complete laws of physics include, besides the dynamical microphysical laws (assumed to be deterministic), a “Past Hypothesis” to the effect that

the universe started in a certain thermodynamic state M_0 of low entropy, as well as a uniform probability distribution `PROB` over microstates realising M_0 . These probabilities, they argue, are best understood as chances in the sense of the best system account. The basic idea is that the package of microphysical laws, Past Hypothesis and `PROB` provides the most elegant and informative summary of the events in our world. The microphysical laws by themselves entail almost nothing about the behaviour of macroscopic objects, unless one happens to know their precise microstate. Adding the probability distribution over initial states, Loewer [2007a: 305] argues, “results in a system that is only a little less simple but is vastly more informative”.

This proposal has been attacked from different directions. Setting aside worries about its physical tenability (see e.g. [Earman 2006]), there are doubts about its application of the best system approach. By the standards of section 3, adding `PROB` and the Past Hypothesis to the micro-laws would not create a system “that is only a little less simple”, since neither addition can be expressed in the fundamental language of microphysics (or if so, only as an infinite disjunction). Loewer therefore remarks that theories may include principles not only in the fundamental language, but also in the language of thermodynamics ([Loewer 2007a: 305, fn.23]). One may wonder what general rule this comes from, and whether the system composed of micro-laws, Past Hypothesis and `PROB` would really come out best. The mere fact that it is better than the micro-laws alone hardly settles that matter. Other worries arise from the fact that Albert and Loewer want the Past Hypothesis to be a *law*, which goes against the usual conception of dynamical laws and non-lawful boundary conditions. (See e.g. Frisch [2011] and Winsberg [2008, 2013] for concerns along these lines.)

Perhaps the most controversial aspect of the Albert-Loewer proposal is the idea that all probabilities in the special sciences are instances of `PROB`. Many philosophers want to grant more autonomy to the special sciences. This is easily achieved in the framework of the best system account: the laws of genetic mutation, for example, can be understood as best systematisations of facts *about genetic mutation*; the basic vocabulary here is not the language of physically fundamental properties, but a suitable language of chemical and biological kinds (see [Schrenk 2007], [Cohen and Callender 2009, 2010]). A variation of this view assumes a single best system for all of science, but allows it to have autonomous parts for different branches of science. The addition of special science laws is then motivated by the fact that they provide systematisations of useful facts about the world that cannot be discerned in the microphysical patterns, or not without impossible computational effort (see [Hoëfer 2007], [Frigg and Hoëfer 2010], [Frisch 2011]).

These applications of the best system account to probabilities in special science demand some modifications to the rules of section 3. Most obvious are issues arising for the specification of the relevant language. In addition, there are different ways of taking into account the fact that most special science laws are only *ceteris paribus* laws. It is

also plausible that new virtues apply to special science theories, such as integration with other areas of science. These details still remain to be spelled out, in what promises to be a fruitful area of future research.

References

- David Albert [2000]: *Time and Chance*. Cambridge (Mass.): Harvard University Press
- Frank Arntzenius and Ned Hall [2003]: “On what we know about chance”. *British Journal for the Philosophy of Science*, 54(2)
- Helen Beebe [2000]: “The Non-Governing Conception of Laws of Nature”. *Philosophy and Phenomenological Research*, 61: 571–594
- Allen R. Bernstein and Frank Wattenberg [1969]: “Non-Standard Measure Theory”. In W. Luxemburg (Ed.) *Applications of Model Theory of Algebra, Analysis, and Probability*, Holt, Reinhart and Winston
- John Bigelow, John Collins and Robert Pargetter [1993]: “The Big Bad Bug: What are the Humean’s Chances?” *British Journal for the Philosophy of Science*, 44: 443–462
- Robert Black [1998]: “Chance, Credence and the Principal Principle”. *British Journal for the Philosophy of Science*, 49: 371–385
- Jonathan Cohen and Craig Callender [2009]: “A Better Best System Account of Lawhood”. *Philosophical Studies*, 145(1): 1–34
- [2010]: “Special Sciences, Conspiracy and the Better Best System Account of Lawhood”. *Erkenntnis*, 73: 427–447
- John Earman [1984]: “Laws of Nature: The Empiricist Challenge”. In R. Bogdan (Ed.) *D. M. Armstrong*, D. Reidel
- [1986]: *A primer on determinism*. Dordrecht: Reidel
- [2006]: “The ‘Past Hypothesis’: Not Even False”. *Studies in History and Philosophy of Science*, 37: 399–430
- John Earman and John T. Roberts [2005]: “Contact with the Nomic: A Challenge for Deniers of Humean Supervenience about Laws of Nature Part I: Humean Supervenience”. *Philosophy and Phenomenological Research*, 71: 1–22
- Adam Elga [2004]: “Infinitesimal Chances and the Laws of Nature”. In Frank Jackson and Graham Priest (Eds.) *Lewisian Themes: The Philosophy of David K. Lewis*, Oxford: Oxford University Press, 68–77

- Roman Frigg and Carl Hoefer [2010]: “Determinism and Chance From a Humean Perspective”. In Friedrich Stadler et al. (Ed.) *The Present Situation in the Philosophy of Science*, Springer, 351–72
- Mathias Frisch [2011]: “From Arbuthnot to Boltzmann: The Past Hypothesis, the Best System, and the Special Sciences”. *Philosophy of Science*, 78: 1001–1011
- Haim Gaifman and Marc Snir [1982]: “Probabilities Over Rich Languages, Testing and Randomness”. *Journal of Symbolic Logic*, 47(3): 495–548
- G. Ghirardi, A. Rimini and T. Weber [1986]: “Unified dynamics for micro and macro systems”. *Physical Review D*, 34: 470–491
- Ronald N. Giere [1973]: “Objective Single Case Probabilities and the Foundations of Statistic”. In P. Suppes et al (Ed.) *Logic, Methodology and the Philosophy of Science IV*, Amsterdam: North Holland, 467–483
- Alan Hájek [1997]: ““Mises redux” – redux: Fifteen arguments against finite frequentism”. *Erkenntnis*, 45(2-3): 209–227
- Ned Hall [1994]: “Correcting the Guide to Objective Chance”. *Mind*, 103: 505–517
- [2004]: “Two Mistakes about Credence and Chance”. *Australasian Journal of Philosophy*, 82: 93–111
- Carl Hoefer [2007]: “The Third Way on Objective Probability: A Skeptic’s Guide to Objective Chance”. *Mind*, 116: 549–596
- James Joyce [2009]: “Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief”. In F. Huber and C. Schmidt-Petri (Eds.) *Degrees of Belief*, Berlin: Springer, 263–300
- Saul Kripke [1980]: *Naming and Necessity*. Oxford: Blackwell
- Marc Lange [2000]: *Natural Laws in Scientific Practice*. Oxford: Oxford University Press
- Isaac Levi [1990]: “Chance”. *Philosophical Topics*, 18: 117–149
- David Lewis [1973]: *Counterfactuals*. Oxford: Blackwell
- [1980]: “A Subjectivist’s Guide to Objective Chance”. In Richard Jeffrey (Ed.), *Studies in Inductive Logic and Probability* Vol. 2, University of California Press, Berkeley.
- [1983]: “New Work for a Theory of Universals”. *Australasian Journal of Philosophy*, 61: 343–377. Reprinted in [Lewis 1999]

- [1986a]: *On the Plurality of Worlds*. Malden (Mass.): Blackwell
- [1986b]: *Philosophical Papers II*. New York, Oxford: Oxford University Press
- [1994]: “Humean Supervenience Debugged”. *Mind*, 103: 473–490
- [1999]: *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press
- Barry Loewer [1997]: “Humean Supervenience”. *Philosophical Topics*, 24: 101–126
- [2001]: “Determinism and chance”. *Studies in History and Philosophy of Science Part B*, 32(4): 609–620
- [2004]: “David Lewis’s Humean Theory of Objective Chance”. *Philosophy of Science*, 71: 1115–1125
- [2007a]: “Counterfactuals and the Second Law”. In H. Price and R. Corry (Eds.) *Causation, Physics, and the Constitution of Reality: Russell’s Republic Revisited*, New York: Oxford University Press, 293–326
- [2007b]: “Laws and Natural Properties”. *Philosophical Topics*, 35: 313–328
- [2012]: “Two accounts of laws and time”. *Philosophical Studies*, 160: 115–137
- David H. Mellor [1971]: *The Matter of Chance*. Cambridge: Cambridge University Press
- John Stuart Mill [1843]: *A System of Logic*. London: Parker
- In Jae Myung, Vijay Balasubramanian and Mark A. Pitt [2000]: “Counting probability distributions: Differential geometry and model selection”. *Proceedings of the National Academy of Sciences*, 97(21): 11170–11175
- P. Pearle [1989]: “Combining stochastic dynamical state-vector reduction with spontaneous localization”. *Physical Review*, 39A: 2277–2289
- Karl Popper [1982]: *Quantum theory and the schism in physics*. Totowa, NJ: Rowman and Littlefield
- Frank P. Ramsey [1978]: “Universals of Law and of Fact”. In *Foundations*, London: Routledge & Kegan Paul
- Jorma Rissanen [1978]: “Modeling by the shortest data description”. *Automatic*, 14: 465–471
- [1989]: *Stochastic Complexity in Statistical Inquiry Theory*. River Edge (NJ): World Scientific Publishing Co.

- Jonathan Schaffer [2007]: “Deterministic Chance?” *British Journal for the Philosophy of Science*, 58: 113–140
- Markus A. Schrenk [2007]: “A Lewisian Theory for Special Science Laws”. In Helen Bohse and Sven Walter (Eds.) *Philosophie: Grundlagen und Anwendungen. Ausgewählte Beiträge aus den Sektionen der GAP 6*, Paderborn: Mentis
- Wolfgang Schwarz [2014]: “Proving the Principal Principle”. In A. Wilson (Ed.) *Chance and Temporal Asymmetry*, Oxford: Oxford University Press, 81–99
- Michael Strevens [1999]: “Objective Probability as a Guide to the World”. *Philosophical Studies*, 95: 243–275
- Michael Tooley [1987]: *Causation: A Realist Approach*. Oxford: Oxford University Press
- Bas C. van Fraassen [1989]: *Laws and Symmetry*. Oxford: Clarendon Press
- Eric Winsberg [2008]: “Laws and chances in statistical mechanics”. *Studies in History and Philosophy of Modern Physics*, 39: 872–888
- [2013]: “The Metaphysical Foundations of Statistical Mechanics: on the status of PROB and PH”. In B. Loewer, E. Winsberg and B. Weslake (Eds.) [*Currently untitled*], Harvard University Press