Modal Epistemology without Modal Facts

Wolfgang Schwarz
Draft, 15 April 2017

1 Introduction

Much recent work in modal epistemology assumes a kind of modal realism according to which reality includes basic modal elements—basic capacities, essences, counterfactuals, etc., which are simply out there, waiting to be discovered. Alternative views of modality put modal epistemology in a very different light. On the reductionist Humean view championed by Lewis (e.g. [Lewis 1986b], [Lewis 1986a], [Lewis 1994]), modal statements express ultimately non-modal propositions concerning the spatiotemporal distribution of categorical properties, and thus modal knowledge is not knowledge of special modal facts. On conventionalist accounts (like [Ayer 1936] or [Sidelle 1989]), modal knowledge is presumably knowledge of linguistic conventions. On projectivist accounts (like [Skyrms 1980] or [Blackburn 1986]), modal knowledge is not knowledge of genuine objective facts at all.

In this essay, I want to explore modal epistemology from a broadly projectivist perspective. The kind of projectivism I have in mind is distinguished from most accounts in the literature by focusing on modal belief rather than modal language. The projectivist hypothesis I want to explore is that our cognitive model of the world includes modal elements whose function is not straightforwardly representational. I will argue that this hypothesis can explain and justify our methods for acquiring modal beliefs a lot better than the realist account commonly assumed in modal epistemology.

My topic is *objective modality*, as opposed to epistemic or deontic modality. That my desk could have been further to the left is a claim of objective modality, as is the claim that my desk would be further to the left if I had pushed it, or that one couldn't get it through my office door without tilting. The outer frontier of objective modality is "metaphysical modality", but to get a clear grasp of metaphysical modality, I think it helps to see it as part of a larger family of objective modality. Here I agree with [Vetter 2015], [Williamson 2016], and others, although for somewhat different reasons. I will turn to metaphysical modality (and these reasons) in section 5. Until then, I will dwell on more restricted kinds of objective modality.

I will begin with a modality whose epistemology has been more thoroughly investigated than that of all the others combined—although the relevant work has mostly been done in statistics and confirmation theory rather than under the heading of modal epistemology. The modality I have in mind is objective probability.

2 Epistemology of objective probability

When thinking about objective probability, it is useful to think in terms of *probability models*. A probability model comprises a "sample space" of possible "outcomes" and a probability measure over that space (or over a measurable sub-algebra of the space, but we may ignore this complication). For example, a probability model for coin flips might consist of the sample space $\{Heads, Tails\}$ and a probability measure that assigns 1/2 to both outcomes. Let's call this model M.

We can't directly observe probabilities. So how are probability models confirmed or disconfirmed by evidence? To answer this question, the framework of Bayesian epistemology has proved useful (see e.g. [Earman 1992], [Howson and Urbach 1993]). It is especially useful for the purposes of the present essay, as it treats questions about confirmation as questions about what it is reasonable to belief, and reasonable belief is our topic.

In Bayesian epistemology, we assume that beliefs come in degrees, and that rational degrees of belief (or credences) satisfy the axioms of probability theory. Since the effect some piece of evidence has on an agent's beliefs generally depends on the agent's background beliefs, it is often convenient to restrict attention to "initial" credence functions: credence functions of hypothetical agents who do not already have relevant information about the propositions at issue.

So let Cr_0 be an arbitrary rational initial credence function. To compute the impact of some evidence E on $\mathrm{Cr}_0(M)$, we often use Bayes' Theorem:

$$\mathrm{Cr}_0(M/E) = \frac{\mathrm{Cr}_0(E/M)}{\mathrm{Cr}_0(E)}\mathrm{Cr}_0(M).$$

Here, $\mathrm{Cr}_0(M/E)$ is the hypothetical agent's degree of belief in the model M conditional on the evidence E, which is also the agent's degree of belief in the model upon learning the evidence. Bayes' Theorem says that $\mathrm{Cr}_0(M/E)$ can be factorised into three parts: the "likelihood" $\mathrm{Cr}_0(E/M)$, the "prior" $\mathrm{Cr}_0(M)$, and the "surprisingness" of the evidence, the reciprocal of $\mathrm{Cr}_0(E)$.

For probability models M, the likelihood is given by another formula, the *Principal Principle*:

$$\mathrm{Cr}_0(E/M \wedge B) = P_M(E).$$

Here, $P_M(E)$ is the probability the model M assigns to E, and B is any "admissible" proposition. What exactly admissibility amounts to is a little controversial and arguably depends on the model, but easily knowable facts about the past of the observation E are generally admissible.

The Principal Principle is really two principles folded into one. The first might be called the *Coordination Principle*. It states that conditional on a probability model, one's (rational initial) credence in relevant outcomes should match the model's probabilities:

$$\mathrm{Cr}_0(O/M) = P_M(O).$$

The second part of the Principal Principle might be called the *Resiliency Principle*, after [Skyrms 1980]. It says that probability models screen off a wide range of other propositions: for any "admissible" B,

$$\mathrm{Cr}_0(E/M \wedge B) = \mathrm{Cr}_0(E/M).$$

To see all this machinery in action, suppose a rational initial credence function is divided between two probability models concerning some kind of coin flip. The first model (M) assigns probability 1/2 to heads, the second (M') probability 2/3. We can now compute how these credences are affected by an observation of a single outcome *Heads* (H). By Bayes' Theorem,

$$\mathrm{Cr}_0(M/H) = \frac{\mathrm{Cr}_0(H/M)}{\mathrm{Cr}_0(H)}\mathrm{Cr}_0(M).$$

By the Principal Principle, $\mathrm{Cr}_0(H/M) = P_M(H) = 1/2$. Moreover, by the law of total probability, $\mathrm{Cr}_0(H) = \mathrm{Cr}_0(H/M)\mathrm{Cr}_0(H) + \mathrm{Cr}_0(H/M')\mathrm{Cr}_0(M')$. Again, by the Principal Principle, this works out to $1/2 \times \mathrm{Cr}_0(M) + 2/3 \times \mathrm{Cr}_0(M')$. Assuming $\mathrm{Cr}_0(M) = \mathrm{Cr}_0(M') = 1/2$, we get

$$\mathrm{Cr}_0(M/H) = \frac{1/2}{1/4 + 1/3}1/2 = 3/7 \approx 0.43.$$

So the hypothetical agent's credence in M will go down from 0.5 to around 0.43. By the same kind of reasoning (and assuming previous outcomes are admissible) we can compute that after observing 5 heads and 5 tails, in any order, the agent's credence in M will have increased to around 0.64; after 21 heads and 20 tails it is 0.89.

In general, if E is information about a sequence of outcomes, then the likelihood $Cr_0(E/M)$ is greater the closer P_M matches the relative frequencies in the sequence. And the greater the likelihood, the greater—all else equal—the boost to the credence in the relevant model. This is how frequencies provide evidence about probability models.

The Principal Principle is not a part of standard probability theory; it is a further norm on rational credence. And it is not the only such norm.

Consider another hypothesis M^*, on which the probability of heads is 2/3 for the first five flips, then 1/2 for the next seven, and from then on 0. Observing $HHTHTTTHTH$ (say) clearly does not rule out M^*; in fact, this sequence gives slightly greater boost to M^* than to M. Nonetheless, rational agents should not give significant credence to gerrymandered hypotheses like M^*. In the Bayesian framework, this can only be avoided by assuming that such hypotheses have negligible prior credence. So we need a further norm on priors: a *Bias for Simplicity and Strength* which implies that $\mathrm{Cr}_0(M^*)$ should be much lower than $\mathrm{Cr}_0(M)$.

The following, more subtle, application of the Bias for Simplicity and Strength will become important later. Let S_1 be a long, disorderly sequence of heads and tails in which both outcomes have equal frequency. Let S_2 be an equally long, but very orderly sequence in which heads is always followed by tails and tails by heads: $HTHTHTHT\ldots$. Intuitively, observing S_1 makes it reasonable to believe that the coin flips are stochastic and fair; observing S_2 does not. But why? Note that the likelihood is the same in either case: conditional on M, S_2 is just as probable as S_1. The reason why S_2 disconfirms M is that it strongly confirms an alternative hypothesis $M^\dagger$ on which the flips deterministically alternate between heads and tails. $M^\dagger$ is reasonably simple and strong and thus—by the Bias for Simplicity and Strength—deserves non-trivial prior credence. Consequently, the posterior credence of $M^\dagger$ after observing S_2 is high, at the expense of all competing hypotheses, including M.

I have not said what exactly the Bias for Simplicity and Strength requires. That's because I don't know. I suspect the norm is not very exact to begin with, but others are more optimistic (e.g. [Solomonoff 1964]). I also suspect that we might need further norms besides the Principal Principle and the Bias for Simplicity and Strength, to explain how information about the physical symmetry of the coin and about the dynamics of the flips supports M. But again I am not sure what exactly these norms should say. For what follows, let's pretend that the Principal Principle and the Bias for Simplicity and Strength are all we need. (As will become clear, adding further norms would only strengthen the case for projectivism.)

3 The cognitive function of probability models

In the previous section, I have reviewed some norms that figure in the confirmation and disconfirmation of probability models by statistical evidence. These norms explicitly constrain rational attitudes towards probability models, but they also put constraints on entirely non-modal beliefs.

For example, suppose you have observed a long and disorderly series of heads and tails, in which both outcomes were about equally common. Absent other relevant information, you should then assign high credence to probability models like M—in accordance with the Principal Principle and the Bias for Simplicity and Strength. Again by the Principal Principle, this implies that you should assign high credence to the hypothesis that the relative frequency of heads in future trials will still be around $1/2$. Thus your non-modal evidence about past outcomes constrains your non-modal beliefs about future outcomes.

The example generalizes. Together, the Principal Principle and the Bias for Simplicity and Strength entail that Cr_0 assigns low credence to worlds where the relative frequency of outcomes suddenly changes. More generally still, they entail that Cr_0 favours worlds with simple, systematic patterns in the history of occurrent events. We may call this a *Bias for Regularity*.

The Bias for Regularity is an independently plausible constraint on rational initial credence: much of our scientific and non-scientific thinking relies on it.

Arguably, the Bias for Regularity makes beliefs about objective probability redundant in the sense that whatever non-modal conclusions you could rationally draw from your total (non-modal) evidence with the help of probability models you could have drawn directly by following the Bias for Regularity. More precisely, if an agent's rational credence is defined over probability models, and she arrives at some credence x in some non-probabilistic proposition A based on (non-probabilistic) evidence E *via* the Principal Principle and the Bias for Simplicity and Strength, then the Bias for Regularity alone would have been enough to license assigning credence x to A (given E).

I have no conclusive proof of this redundancy claim, as that would require a precise formulation the relevant norms. For what it's worth, here is an outline of an argument:

> Let Cr_0 be a rational (initial) credence function whose domain includes hypotheses about objective probability; let Cr_0^H be a rational "Humean" credence function defined only over non-modal propositions. Cr_0 satisfies the Principal Principle and the Bias for Simplicity and Strength; Cr_0^H instead satisfies the Bias for Regularity. For any probability model M, let M^H be the hypothesis that M best combines the virtues of simplicity, strength, fit (as in [Lewis 1994]). Let PP^H and BSS^H be the corresponding transformations of the Principal Principle and the Bias for Simplicity and Strength; i.e., PP^H says that for admissible B, $Cr_0(E/M^H \wedge B) = P_M(E)$, and BSS^H says that $Cr_0(M^H)$ is greater the simpler and stronger M. Plausibly, the Bias for Regularity entails BSS^H. In [Schwarz 2014], I've argued that (together with independently plausible principles of rationality) it also entails PP^H. Assume for simplicity that $\mathrm{Cr}_0(M/E)$ is fully determined by the Principal Principle and the Bias for Simplicity and Strength. (The argument can be generalised to more permissive accounts.) It follows that $\mathrm{Cr}_0(M/E) = \mathrm{Cr}_0^H(M^H/E)$. Moreover, suppose A is in the domain of the probability model M and E is admissible for A (relative to M and Cr_0). Then $\mathrm{Cr}_0(A/M \wedge E) = P_M(E) = \mathrm{Cr}_0^H(A/M^H \wedge E)$, because Cr_0^H satisfies PP^H. Now if $\{M_i\}$ is some partition of probability models, then by the law of total probability, $Cr_0(A/E) = \sum_i Cr_0(A/M_i \wedge E)Cr_0(M_i/E)$ and $Cr_0^H(A/E) = \sum_i Cr_0^H(A/M_i^H \wedge E)Cr_0(M_i^H/E)$. As we've just shown, $Cr_0(A/M_i \wedge E) = Cr_0^H(A/M_i^H \wedge E)$ and $Cr_0(M_i/E) = Cr_0(M_i^H/E)$. So $\mathrm{Cr}_0(A/E) = \mathrm{Cr}_0^H(A/E)$. QED.

The redundancy claim is trivial if we stipulatively define the Bias for Regularity as whatever constraint the norms governing attitudes towards probability models impose on non-probabilistic beliefs. But that's not how I want to define it, for that would leave open whether the Bias is plausible as a basic Humean norm. The somewhat non-trivial claim I want to put forward is that there are independently plausible Humean norms on rational credence (such as the Bias for Regularity) the endorsement of which makes beliefs about objective probability redundant.

I speak of 'Humean' norms because anti-Humeans sometimes hold that it is only rational to believe that the world is regular if one believes that its regularity is grounded in non-Humean powers, laws of nature, chances, or the like. On that view, the Bias for Regularity is not a basic norm of rationality. It is derived from other norms, some of which (like the Bias for Simplicity and Strength) require giving high credence to worlds where the non-Humean powers, laws, chances, etc. are reasonably simple and constant. Such norms are still needed: merely postulating non-Humean laws, for example, does not guarantee that the behaviour of physical systems in the world will be orderly, unless the relevant laws are reasonably simple and constant. Presumably, the assumption that the laws (or powers etc.) are orderly cannot be justified by anything more fundamental. Thus on the anti-Humean view, the orderliness of the world is explained by laws, but the orderliness of the laws is not explained by anything else; it is a brute fact that simply has to be believed, even in the absence of relevant evidence.

So everyone has a bump in their rug. Everyone needs a primitive bias for orderliness, it's just that Humeans and anti-Humeans locate it at different places. However, the anti-Humean rug has further bumps, and larger bumps, for she has to postulate additional primitive norms (like the Principal Principle), which the Humean doesn't need.

I will turn to this matter in the next section. First let me return to the redundancy claim and what it suggests about the cognitive role of probability models. If the claim is true, beliefs about objective probability are not needed to guide ordinary, non-modal beliefs. Instead of endowing us with a capacity to have beliefs about objective probability and ensuring that these beliefs approximately satisfy the Principal Principle, the Bias for Simplicity and Strength, and possibly other norms, evolution could have restricted our doxastic space to ordinary propositions and instead made sure we satisfy the Bias for Regularity. Why did evolution choose the first route?

One possible answer is that beliefs about objective probability are useful not just because of the way they guide non-probabilistic beliefs. What could the other uses be? The most promising answer I can think of is that beliefs about objective probability are directly useful for guiding actions: on Lewis's ([Lewis 1981]) and Skyrms's ([Skyrms 1980], [Skyrms 1984]) formulation of causal decision theory, the definition of expected utility involves the agent's credence in hypotheses about objective probability. However, it is doubtful that these credences are indispensable and irreducible. Indeed, Lewis and Skyrms are both Humeans. For Lewis, the relevant hypotheses about objective probability are really hypotheses about certain patterns in the distribution of categorical properties. For Skyrms, what appear to be credences over probability models are really just credences over non-modal propositions that satisfy certain invariance conditions.

I want to suggest a different explanation of why evolution choose the apparently more complex option of giving us credences about objective probability. The explanation is that this actually reduces our cognitive load: probability models are highly efficient tools for storing and manipulating complex (non-modal) information.

Think of the difficulty of computing the trajectory of all molecules in a container of gas. The task is utterly intractable, even with a very small number of molecules. By comparison, statistical models that abstract away from the precise position and momentum of individual molecules are almost trivial to use, and their probabilistic predictions often suffice to answer our questions. What's more, their answers are often more useful and perspicuous than the answers we could have gotten from a detailed study of the lower-level dynamics.

The same is true for probability models in other domains. There is a simple probabilistic explanation of why the genetic diversity of isolated, small populations generally decreases—an explanation that would be hard to see if you traced the individual alleles in any given population over time.

In artificial intelligence, hardly anyone even considers agents with a single, "joint" credence function over all relevant ways the world could be. The agent's credences are almost always broken down into a number of components, including simple probability models that specify (often conditional) probabilities over a highly restricted domain; these models can themselves be learned, and sophisticated agents can have subjective probabilities over the probability models. Evolution may well have used the same trick when it designed us.

4 The case for projectivism

If beliefs about objective probability are cognitive tools to store and manipulate complex non-modal information, it would be surprising if they simultaneously managed to track irreducible probability facts in the world. Those who believe in such facts face the challenge to explain why evolution should have endowed us with a capacity to recognise and reason about these facts, especially if any non-probabilistic conclusions we can draw with the help of the capacity could in principle also be drawn without it.

The projectivist account, on which beliefs about objective probability are *mere* cognitive tools, also fares a lot better than realist accounts when it comes to explaining and vindicating our epistemic practice: our methods of assessing and testing hypotheses about objective probability.

As we saw, a central aspect of our epistemic practice is encoded in the Principal Principle. As Lewis famously complained, taking objective probability (or "chance") as primitive would make a mystery out of that Principle:

> I haven't the faintest notion how it might be rational to conform my credences about outcomes to my credences about some mysterious unHumean magnitude.

> Don't try to take the mystery away by saying that this unHumean magnitude is none other than *chance*! I say that I haven't the faintest notion how an unHumean magnitude can possibly do what it must do to deserve that name. . . . [Lewis 1986b: xvf.]

Lewis's complaint is sometimes met with bewilderment (see e.g. [Hall 2004]), but it makes an important point. Suppose there is some metaphysically primitive probabilistic quantity X. Then it is presumably possible for somebody to not align their degrees of belief (as reflected, for instance, in their actions) with what they believe about that quantity. If objective probability is identified with X, the Principal Principle therefore turns into the substantive requirement that one's credences should be aligned with one's beliefs about X. The requirement sounds harmless if we refer to X as 'objective probability' or 'chance'. On Lewis's view, the requirement then actually becomes analytic, for 'chance' can be defined as whatever quantity satisfies the Principal Principle. Proposing that chance is a metaphysically primitive quantity therefore *is* to propose that one should align one's credences with one's beliefs about a metaphysically primitive quantity. And that is a substantive claim, not a trivial tautology. Indeed, the claim is surely a little perplexing. One would like to know why rationality requires aligning one's credences about outcomes with one's beliefs about a primitive quantity that is (metaphysically) independent of the outcomes. Nobody has ever offered an explanation. Anti-Huemans typically agree that the requirement is a basic norm of rationality.

An analogy with morality might help. Suppose for the sake of the analogy that a common method for acquiring moral beliefs is to consider what one would desire under certain circumstances. This would seem to present a challenge to moral realists who hold that moral facts are a metaphysically basic dimension of reality and therefore (metaphysically) independent of what people desire or what they would desire under various circumstances: why should introspecting one's state of desire provide evidence about an independent and primitive aspect of reality? The realist might respond that it is simply a basic norm of rationality that we should align our beliefs about what we would desire with our beliefs about what is right and wrong. But this is hardly a satisfactory answer to the challenge. By postulating a basic norm, the moral realist would effectively concede that no explanation can be given for our access to moral facts. Similarly, realists about probability effectively concede that our access to probability facts is an inexplicable mystery.

To be sure, at some point explanations—in epistemology as elsewhere—must come to an end. Perhaps not much more can be said to explain why it is rational to infer that unobserved emeralds are green from the fact that observed emeralds are green, even though these propositions are metaphysically independent. But postulating primitive norms of rationality should be a last resort. At the very least, it should count as a cost for non-reductive realists (both about modality and morality) that they have to postulate new norms.

Remember also that the Principal Principle is really two norms: Coordination and Resiliency. If objective probability is metaphysically primitive, both of these presumably have to be accepted as basic norms of rationality. While Coordination has drawn the most attention, the claim is perhaps even harder to swallow for Resiliency, which involves a somewhat peculiar restriction to admissible propositions. It is hard to accept that nothing can be said to explain why the principle is restricted in just that way.

On the projectivist account I have outlined, objective probability is not identified with any objective quantity (reducible or irreducible), so the Principal Principle is not a substantive norm on one's beliefs about the world. Rather, it is a rule for using probability models. The need for a rule like Coordination is obvious: extending one's doxastic space by new elements would be pointless if there weren't rules connecting credences over the new elements with credences over the original elements. Resiliency is important because it allows the decomposition of joint probabilities without which probability models would lose much of their advantage for cognitive economy. It is also not surprising that information we can easily come to have (in advance of the relevant event) is generally admissible: again, that is required to make probability models useful in practical applications.

Don't we still have to postulate that agents *should* use probability models in accordance with the Principal Principle? Perhaps. But if so, that is not a basic norm. For recall that the proper use of probability models is in principle redundant. If anyone improperly used the models by violating the Principal Principle, they would also violate the Bias for Regularity or some other Humean norm of rationality.

So the projectivist account neatly explains and vindicates our methods for acquiring beliefs about objective probability. On these terms, no extant realist account comes even close.

I have focused on non-reductive forms of realism. But there are also reductive accounts of objective probability, the most popular of which is the "Best-Systems Account" (defended e.g. in [Lewis 1994] and [Loewer 2004]) on which statements about objective probability describe complex regularities in the history of occurrent events. Epistemically, I think the projectivist account and the Best-Systems Account are roughly on a par: both can explain our epistemic practice, neither needs extra norms of rationality. However, the Best-Systems Account clashes with a wide range of intuitions that might be considered central to our conception of objective probability (see e.g. [Bigelow et al. 1993]). For example, on the Best-Systems Account, present probabilities depend on future outcomes: if the all radium atoms from now on were to decay within 100 years, the present half-life of radium would likely be different. Relatedly, it seems that one can clearly imagine scenarios with the very same history of outcomes but different objective probabilities (see e.g. [Tooley 1987: 42–47]). Moreover, it is arguably a common practice of science to study simple scenarios governed by a given system of laws and chances even though on the Best-Systems Account the scenario cannot have those laws and chances (see e.g. [Maudlin 2007: 67f.]). The projectivist account avoids all these costs because it treats probability hypotheses as independent elements in our doxastic space. We can therefore hold these elements fixed when studying toy scenarios or when we evaluate counterfactuals about the future decay of atoms. We can also accept that any given hypothesis about the total history of occurrent events is compatible with a wide range of hypotheses about objective probability.

The projectivist account thus avoids a dilemma that arises for realist accounts: on the one hand, objective probability seems to be independent of occurrent events; but if it is, then it becomes mysterious how observing occurrent events licenses inferences about objective probability.

The most serious cost of projectivist is perhaps its revisionary character: that it denies the reality of objective probability. Other notorious problems for projectivism are avoided by the shift from language to belief. For example, the "Frege-Geach problem" arguably does not arise for the present form of projectivism. If our doxastic space includes probability models, assertions about objective probability straightforwardly express beliefs, and there is no mystery about what might be expressed (say) by conjunctions of such statements with other statements. Nor is there a problem about specifying truth-conditions, provided the truth-conditions can be specified in terms of the things speakers and hearers believe in. There may be complications when we try to specify truth-conditions in terms of what really and fundamentally exists. But why should we try to do that? To explain facts about communication or linguistic processing, truth-conditions of the first kind are much more suitable.

5 Beyond probability

I have talked a lot about objective probability. What about other forms of objective modality?

Our concept of objective probability belongs to the family of "nomic" concepts. Other members of that family include physical possibility and necessity, causality, and certain counterfactuals. There are close conceptual connections between these concepts (in many directions), so a uniform treatment seems desirable.

In fact, we have already seen one use for beliefs about nomic necessity. Recall the scenario from page 4 where you observe a coin alternating between heads and tails. To explain why this disconfirms the stochastic hypothesis M on which the probability of heads on each toss is $1/2$, we had to assume that there is an alternative $M^\dagger$ to M on which the coin deterministically alternates between the two outcomes. For the explanation to go through, $M^\dagger$ cannot be just a description of the alternating outcomes, as that would be compatible with M. Rather, $M^\dagger$ has to say that heads *must* be followed by tails, in a sense in which that is incompatible with the outcomes being a matter of probability. The lesson illustrated by this example is that in order to fulfil its cognitive function, the space of probability models has to include deterministic models.

Instead of exploring further how the projectivist approach I have outlined might be generalized to other nomic elements, I want to conclude by briefly turning to metaphysical modality. Here a special treatment seems to be called for.

Beliefs about objective probability are useful tools for storing and manipulating non-modal information. The same can arguably be said for beliefs about nomic possibility and necessity, but not for beliefs about metaphysical modality. For example, I believe that there could have been no limit to how fast things can travel. That belief does not allow me to predict anything, nor does it seem to reflect any empirical information I have about the world. So what is the point of the belief? Or take the common belief (among philosophers) that Kripke could not have been a scrambled egg. This belief arguably does reflect empirical information—that Kripke is human. But it doesn't provide an efficient summary of such information, and again it doesn't seem to license any interesting predictions. Why do we have a capacity for such beliefs?

Realism about metaphysical modality does not offer an explanation. Let's grant that there are basic truths about what's metaphysically necessary and possible. Why should we be interested in these truths? Why should evolution have given us a capacity to track them? Some hold that Kripke actually could have been a scrambled egg. If that is false, it does not seem to cause any difficulties for those who hold the belief. How could false beliefs like this have impaired the evolutionary fitness of our ancestors?

It is evidently useful to have a capacity for reasoning about what is possible in a sense that goes beyond nomic possibility—for example, when we consider different hypotheses about the laws of nature as possible explanations of a given phenomenon. But such hypotheses are best understood as *epistemic* possibilities. They are not constrained by the kind of things we tend to hold fixed when considering what is metaphysically possible. For example, we can easily entertain different hypotheses about the chemical structure of gold as possible explanations of certain phenomena, even if we think that gold has its chemical structure essentially.

In a Bayesian (or Hintikkaen) framework, epistemic modality comes for free: any proposition that is not ruled out by the agent's evidence is an epistemic possibility. We can also define a notion of *deep* epistemic possibility or apriority: any non-empty proposition in the agent's probability space is deeply epistemically possible. To be sure, these are merely formal definitions. There are difficult and important questions about the nature and extent of an agent's doxastic space; the present essay deals with a small sub-problem of this question. My point is only that Bayesian (or Hinikkaen) agents don't need special models of epistemic modality.

It has often been observed that there are interesting connections between apriority and metaphysical necessity. Specifically, in every case (or, at any rate, in every clear case) where a claim is necessary but not a priori, the necessity seems to follow a priori from ordinary non-modal facts about the world. Any account of metaphysical modality had better explain this observation.

One attractive explanation is offered by the "two-dimensional" account of modality, which draws on another observation from the semantics of context-dependent expressions. Consider a statement like 'I am not here'. There is a sense in which the statement is false in every possible situation—perhaps bracketing some strange and unusual possibilities. Nonetheless, one can truly say 'if I had gone to Paris, I would not be here', and it is tempting to analyse the conditional by evaluating the consequent 'I am not here' in counterfactual situations that satisfy the antecedent 'I am in Paris'. For this to work out, 'I am not here' must evaluate as true in the relevant situations where I am in Paris. But there may be nothing strange or unusual about these. So there seem to be two ways of evaluating a sentence at a possible situation: 'I am not here' is false in every ordinary situation *considered as actual*, but true in many situations *considered as counterfactual*.

According to the two-dimensional analysis of modality (defended, in different flavours, by Lewis [1986a], Jackson [1998], Chalmers [2006], and others), a statement is metaphysically necessary just in case it is true in all possible situations considered as counterfactual, while it is a priori if it is true in all situations considered as actual. The situations are the same in either case.

So once we have a space of deep epistemic possibility, we may not need a further space of metaphysical possibility to account for beliefs about metaphysical modality. Return to the claim that Kripke could not have been a scrambled egg. On the two-dimensional analysis as developed for example by Lewis, the claim is true (evaluated as actual) at a (deeply) epistemically possible world w just in case whoever stands at the relevant origin of our use of the name 'Kripke' at w does not have a scrambled egg as a counterpart at any other (deeply) epistemically possible world. Equivalently, the claim is true at w just in case whoever stands at the relevant origin of our use of the name 'Kripke' at w is dissimilar (in a specific sense) to a scrambled egg. That explains why the claim is a posteriori, and also why the information you need to figure out that Kripke could not have been a scrambled egg is ordinary non-modal information: what you need to know is that 'Kripke' picks out an object that's sufficiently dissimilar to a scrambled egg.

But a question remains. What's the point of evaluating expressions at situations considered as counterfactual? If all you want to say is that Kripke is dissimilar to a scrambled egg, why would you express that by 'Kripke could not have been a scrambled egg'?

This is where I suspect more restricted forms of objective modality come into play.

Suppose you accept a probabilistic model on which some genetic defect increases both an agent's desire to smoke and her propensity to develop cancer, which fully explains the known correlation between smoking and cancer. You wonder whether to start smoking. The hypothesis that you will smoke raises your credence in having the gene defect and thus in getting cancer. But arguably this is not a good reason not to smoke. To properly evaluate whether you should smoke, you have to consider the relevant smoking situations 'as counterfactual' rather than 'as actual': you have to consider them relative to your (partly unknown) actual situation. If in fact you don't have the gene defect, you still wouldn't have it if you started smoking; if in fact you do have the gene defect, you would still have it if you started smoking.

In general, applications of probabilistic models often involve assessing the probability of certain events under conditions that are explicitly regarded as counterfactual. It is no surprise that we have linguistic rules that facilitate expressing these kinds of applications. When we consider a situation as counterfactual, we typically want to relate it in some way to the actual situation. So it makes sense to have an expression like 'here' that picks out our actual location even in situations considered as counterfactual in which we are somewhere else. And what's true for locations is true for many other things: it will be convenient to have "rigid" expressions whose denotation in situations considered as counterfactual is some function of their actual denotation.

References

Alfred Jules Ayer [1936]: *Language, truth and logic*. Dover Publications

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

Simon Blackburn [1986]: "Morals and Modals". In G. MacDonald and C. Wright (Eds.) *Fact, Science, and Morality,* Oxford University Press, 119–142

David Chalmers [2006]: "The Foundations of Two-Dimensional Semantics". In Manuel Garcia-Carpintero and Josep Macia (Eds.) *Two-Dimensional Semantics,* Oxford: Oxford University Press, 55–140

John Earman [1992]: *Bayes or Bust?*. Cambridge, MA: MIT Press

Ned Hall [2004]: "Two Mistakes about Credence and Chance". *Australasian Journal of Philosophy*, 82: 93–111

Colin Howson and Peter Urbach [1993]: *Scientific Reasoning*. La Salle (Ill.): Open Court Press, 2nd edition

Frank Jackson [1998]: *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Clarendon Press

David Lewis [1981]: “Causal Decision Theory”. *Australasian Journal of Philosophy*, 59: 5–30

— [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

Barry Loewer [2004]: “David Lewis’s Humean Theory of Objective Chance”. *Philosophy of Science*, 71: 1115–1125

Tim Maudlin [2007]: *The metaphysics within physics*. Oxford: Oxford University Press

Wolfgang Schwarz [2014]: “Proving the Principal Principle”. In A. Wilson (Ed.) *Chance and Temporal Asymmetry,* Oxford: Oxford University Press, 81–99

Alan Sidelle [1989]: *Necessity, Essence, and Individuation: A Defense of Conventionalism*. Cornell University Press

Brian Skyrms [1980]: *Causal Necessity. A Pragmatic Investigation of the Necessity of Laws*. New Haven: Yale University Press

— [1984]: *Pragmatics and Empiricism*. Yale: Yale University Press

Ray J Solomonoff [1964]: “A formal theory of inductive inference. Part I”. *Information and control*, 7(1): 1–22

Michael Tooley [1987]: *Causation: A Realist Approach*. Oxford: Oxford University Press

Barbara Vetter [2015]: *Potentiality: from Dispositions to Modality*. Oxford University Press

Timothy Williamson [2016]: “Modal science”. *Canadian Journal of Philosophy*, 46(4-5): 453–492