

De Finetti's Gambit

Wolfgang Schwarz

Draft, February 2025

1 Introduction

Hume [1739] argued that it is a mistake to seek metaphysical explanations for what makes some propositions necessary, or some acts wrong. To understand necessity or morality, Hume argued, we must turn to psychology: we should ask what it is to *judge* that something is necessary or wrong. To say that an act is wrong might simply express a kind of disapproval. To say that a proposition is necessary might express a disposition to hold it fixed in certain kinds of counterfactual reasoning. We don't need wrongness or necessity in our metaphysics.

Followers of Hume have tried to spell out this idea – see, for example, the classical essays in [Blackburn 1993]. The project faces some deep problems. What exactly is the non-doxastic attitude expressed by our judgements, and how can it explain the cognitive role of these judgements? How can we make sense of mixed judgements that appear to be partly about morality or necessity and partly about real aspects of the world? How can we make sense of uncertainty towards moral or modal propositions?

In this essay, I want to look at a particular instance of the projectivist program, dealing with physical probability, or chance. Hume [Hume 1739: bk.I pt.III sec.XI] argued that the concept of chance is a projection of our beliefs onto the world. Bruno de Finetti [1970] gave mathematical substance to Hume's proposal. According to an impressive list of philosophers and statisticians, de Finetti showed that scientific reasoning about chance is best understood as arising from symmetries in degrees of belief: we don't need objective chance in our metaphysics. (See, for example, [Bernardo and Smith 2009], [Spiegelhalter 2024], [Skyrms 1984: ch.3], [Diaconis and Skyrms 2017: ch.7], [Jeffrey 2004].)

Unfortunately, presentations of de Finetti's argument tend to be heavy on the maths and thin on the philosophy. As a result, the projectivist theory of chance has not become widely accepted. Science textbooks generally ignore it. So does the excellent Stanford Encyclopedia entry on "Interpretations of Probability" [Hájek 2023].

My aim in this essay is to fill in the missing philosophical details. I'll explain how de Finetti's gambit is supposed to work, and how it overcomes the problems that beset traditional developments of expressivism and projectivism. The key move that makes this possible is to start not with language and assertion, but with credence.

2 Credence and chance

Subjective probability is well understood. Beliefs come in degrees, and there are good reasons to think that ideally rational degrees of belief should satisfy the rules of probability. I won't review these matters.

Objective chance, on the other hand, is not well understood. There is no agreement on what it means to say that a coin has a 50% chance of landing heads, or that a radium-227 atom has a 50% chance of decaying within 42 seconds.

It's clear that these can't be *directly* understood as concerned with belief. If we're unsure about chance, we're not unsure about our beliefs, or about anyone else's beliefs. A naive subjectivist interpretation of chance would turn almost all of science into a branch of psychology (or worse, of *normative* psychology, if chance is analyzed in terms of *rational* belief). This will not do.

Science textbooks commonly define probability as relative frequency: "the number of favourable cases" divided by "the number of all the cases", as [Laplace 1814] put it. This can't be right either. It would make it logically impossible that a fair coin is tossed an odd number of times. Switching to limiting relative frequencies in hypothetical trials, as some textbooks suggest, hardly make things better. It is doubtful that there is a fact of the matter about how a coin would land if it were tossed infinitely often. Any world where this happens would surely have to be quite unlike ours: why should science be obsessed with such worlds?¹

If chance can't be defined in terms of anything else, perhaps we should take it as fundamental. On this account, the physical universe involves a basic probabilistic quantity. Probability statements in science are simply statements about that quantity.

This view also faces serious problems. For one thing, we find probability statements not just in fundamental physics, but also in statistical mechanics, systems biology, population ecology, epidemiology, and many other fields that are commonly agreed to not require probabilities at the level of fundamental physics. The chances of coin flips and dice rolls don't seem to be physically fundamental. It's therefore hard to see how they could be metaphysically fundamental.

Besides, how could we learn about fundamental chances? Consider what God would have to do to create a world with fundamental chances – what she would have to do to pick out such a world from the set of all possible worlds. One thing she would have to do is settle the chances. Imagine she turns a dial to settle the chance of Heads for a truly indeterministic type of coin flip. In addition, she also has to choose the outcome of each coin flip: that the first lands Heads, the second Tails, and so on. The two choices are almost entirely independent. Unless God has set the dial to 0 or 1, she is free to choose any outcomes she wants. After all, any combination of non-trivial chance hypotheses is compatible with any sequence of outcomes. A coin that's biased towards Heads can still land Tails on each trial. (This is another reason why the relative frequency interpretation is untenable.) If we live in a world that has been created in this manner, how could we get information about how God has set the dial? How could it help to look at the outcomes, or to scrutinize the coins?

The severity of this problem is easy to miss because we know very well how to reason from outcomes to chances. In what follows, I will assume a broadly Bayesian account of scientific reasoning,

¹ See, e.g., [Hájek 1997] and [Hájek 2009] for further arguments against frequentist interpretations.

as defended, for example, in [Howson and Urbach 1993] and [Earman 1992]. The inference from outcomes to chances here works as follows.

We start with a prior credence Cr over chance hypotheses. In the coin example, this might be a credence over the chance θ of Heads on individual flips. Suppose we now observe a sequence of outcomes x_1, \dots, x_n with $\#h$ Heads and $\#t$ Tails. In response, we update our credence towards θ by *Bayes' Rule*:

$$\text{Cr}(\theta \mid x_1, \dots, x_n) \propto \theta^{\#h} (1 - \theta)^{\#t} \times \text{Cr}(\theta). \quad (\text{BR})$$

The larger the observed sequence, the more this concentrates the posterior credence around θ values near the observed frequencies.

Bayes' Rule can be derived from Bayes' Theorem, an independence assumption, and a version of the Principal Principle. Bayes' Theorem says that

$$\text{Cr}(\theta \mid x_1, \dots, x_n) \propto \text{Cr}(x_1, \dots, x_n \mid \theta) \times \text{Cr}(\theta). \quad (\text{BT})$$

The independence assumption is that conditional on a particular chance value θ , no outcome provides evidence about any other outcome:

$$\text{Cr}(x_1, \dots, x_n \mid \theta) = \text{Cr}(x_1 \mid \theta) \times \dots \times \text{Cr}(x_n \mid \theta). \quad (\text{Ind})$$

Finally, the required version of the Principal Principle says that our credence in a Heads outcome conditional on Heads having chance θ equals θ :²

$$\text{Cr}(x_i = H \mid \theta) = \theta. \quad (\text{PP})$$

Bayes' Rule is an immediate consequence of these three assumptions.

To complete the story, we need to explain what the credence function Cr is supposed to represent. On a radical "subjectivist" interpretation, it represents the actual degrees of belief of whoever is engaged in the reasoning. On a radical "objectivist" interpretation, Cr is a fixed confirmation measure to which actual beliefs *ought to conform*. Various intermediate interpretations are possible. We don't need to take a stance on this issue. Let's assume that Cr is the credence function of a rational scientist, leaving open what exactly this demands.

I assume that the Bayesian story captures how we can reason from outcomes to chances. So *how* is not really the problem. The problem is to explain why this method works. If chance is relative frequency, for example, then Bayes' Rule is demonstrably unsound.³ If chance is fundamental, the rule is formally consistent, but its justification becomes a mystery. Why should we be highly confident that God has set the dial to near 1/2 just because we observe an equal number of Heads and Tails? What

² This is not the Principal Principle of [Lewis 1980]. Lewis's Principle assumes that chance is defined for arbitrary propositions. Most probabilities in science arguably don't satisfy this assumption.

³ On the frequency interpretation, each outcome in a finite sequence provides evidence about the others, even conditional on a particular chance hypothesis. So (Ind) becomes untenable. The "Best Systems" interpretation of chance of [Lewis 1994] has the same problem. If we assume that the chance function itself treats the trials as independent, we also get a violation of the Principal Principle. In the literature on the Best Systems account, this is known as the "undermining problem".

do the outcomes and the dial setting have to do with one another, given that both are metaphysically basic and (metaphysically) independent?

This is where de Finetti enters the stage. De Finetti saw that if we understand chance as a projection from our credences then we can explain our methodological practice: we can *derive* Bayes' Rule. Let me explain how this works.

3 De Finetti's Theorem

In what follows, I'll pretend that the world is a sequence of coin flips. This is obviously incorrect, but it helps to keep the mathematics (somewhat) simple.

So assume we have a rational credence Cr over sequences of outcomes. The sequences may have any length. Assume, further, that the credence assigned to a sequence only depends on the number of Heads and Tails, not on their order. For example, *HHTHTTHT* and *THHHTTHT* have the same credence, because they have the same number of Heads and Tails. In this case, we say that the credence is *exchangeable*.

In essence, de Finetti's Theorem now states that this credence can be *represented* as a probability function that

- is unsure about the chance θ on each trial,
- regards the trials as independent, and
- satisfies Bayes' Rule.

To clarify the nature of this "representation", I'll briefly sketch a proof of the Theorem, although most of the details won't be important for what follows.⁴

Assume that x_1, \dots, x_n is a finite sequence of outcomes, drawn from a possibly longer "world" sequence x_1, \dots, x_N . We assume that Cr is exchangeable, and that the probability it assigns to x_1, \dots, x_n does not depend on N : mere information about the size of the world sequence doesn't reveal anything about the initial outcomes.⁵ Choose a particular $N > n$, and let R be a variable for the ratio of heads in x_1, \dots, x_N . By the law of total probability,

$$\text{Cr}(x_1, \dots, x_n) = \sum_r \text{Cr}(x_1, \dots, x_n \mid R = r) \text{Cr}(R = r). \quad (1)$$

By exchangeability, all sequences x_1, \dots, x_N with the same number of heads have the same probability. This means that $\text{Cr}(x_1, \dots, x_n \mid R = r)$ equals the distribution one gets when drawing n balls (without replacement) from an urn with N balls, r of which are white and the others black. The larger N is compared to n , the more this distribution looks like the binomial distribution for drawing without replacement. More precisely, if the sequence x_1, \dots, x_n contains $\#h$ heads, and the "world" sequence x_1, \dots, x_N contains $\#H$ ($= rN$) heads, then there are $\binom{N}{\#H}$ overall possibilities for x_1, \dots, x_N , and

⁴ My proof follows [Diaconis and Freedman 1980] rather than de Finetti, who assumes that Cr is defined over infinite sequences.

⁵ That is, the marginals satisfy Kolmogorov's consistency conditions.

$\binom{N-n}{\#H-\#h}$ possibilities whose initial segment is x_1, \dots, x_n . So

$$\text{Cr}(x_1, \dots, x_n \mid R = r) = \frac{\binom{N-n}{rN-\#h}}{\binom{N}{rN}}. \quad (2)$$

Putting (1) and (2) together, we have

$$\text{Cr}(x_1, \dots, x_n) = \sum_r \binom{N-n}{rN-\#h} / \binom{N}{rN} \text{Cr}(R = r). \quad (3)$$

Because Cr puts no upper bound to the size N of the world, this holds for any $N > n$. As we increase N , the ratio R converges (with probability 1) to a limit θ ; $\binom{N-n}{rN-\#h} / \binom{N}{rN}$ converges to the binomial $\theta^{\#h} (1 - \theta)^{\#t}$; the sum becomes an integral, and $\text{Cr}(R = r)$ turns into a (unique) probability density μ over $\theta \in [0, 1]$. That is, there is a unique μ such that

$$\text{Cr}(x_1, \dots, x_n) = \int_0^1 \theta^{\#h} (1 - \theta)^{\#t} \mu(d\theta). \quad (\text{FT})$$

This is de Finetti's Theorem.

The right-hand of (FT) “represents” a probability function that is unsure about the chance and regards the tosses as independent insofar as μ is unsure about the chance parameter θ , and $\theta^{\#h} (1 - \theta)^{\#t}$ is the probability of the outcomes on the assumption that tosses are independent and that the probability of heads on each toss is θ .

De Finetti's idea is all it *means* to have a credence over chances is to have a corresponding exchangeable credence over outcomes. The “credence over chances” implicit in Cr is given by the (unique) density μ in (FT). Also implicit in Cr are credences about outcomes conditional on chance. These are given by $\theta^{\#h} (1 - \theta)^{\#t}$. If we abuse notation and write ‘Cr(θ)’ for the implicit credence over chances, and ‘Cr($x_1, \dots, x_n \mid \theta$)’ for the implicit credence over outcomes given the chance, Bayes' Theorem gives us an inverse credence ‘Cr($\theta \mid \omega$)’ over chances given outcomes. Specifically:

$$\text{Cr}(\theta \mid x_1, \dots, x_n) \propto \theta^{\#h} \cdot (1 - \theta)^{\#t} \cdot \text{Cr}(\theta).$$

Bayes' Rule!

On de Finetti's interpretation, scientific reasoning, with its apparent commitment to objective chance, can at most be accused of sloppy notation. It is mathematically sound and philosophically innocent, with no real commitment to objective chance. As Diaconis and Skyrms [2017: p.123] put it: “if we dispense with objective chance, nothing is lost. The mathematics of inductive reasoning remains exactly the same.”

To summarize. Realist accounts of chance have trouble explaining the role of chance in scientific reasoning. De Finetti's projectivist re-interpretation explains this role as a mathematical consequence of a certain symmetry (exchangeability) in rational credence functions.

4 Towards quasi-realism

We've seen how de Finetti's re-interpretation of chance can vindicate "the mathematics of inductive reasoning". But physical probability doesn't just figure in the mathematics of inductive reasoning. We also find it, for example, in scientific laws. How should we understand these laws? Do radium-227 atoms not have a half-life? Or is their half-life dependent on our credence, so that they might have one half-life for you and another for me?

A strict de Finettian does not have beliefs about chance. 'Cr(θ)' is just sloppy notation, shorthand for $\mu(\theta)$, where μ is the density in the de Finetti representation of Cr. Accordingly, a strict de Finettian wouldn't say that there is a chance of heads, or that radium atoms have a half-life. Their credence function is only defined over occurrent events, and chance hypotheses are not definable in terms of such events. Chance does not exist.

In its strict form, de Finetti's gambit is a revolutionary form of error theory. It calls for a radical revision of our attitudes. If this is what Diaconis and Skyrms have in mind, it is surely an exaggeration to suggest that "nothing is lost".

I want to explore a less revolutionary version of de Finetti's gambit that allows for genuine beliefs about chance. It treats some statements about chance as true and others as false, and vindicates the idea that the chances do not depend on us. On this interpretation, the gambit approaches a kind of *quasi-realism*, in the sense of Blackburn [1993]. (Indeed, it is evidently related to [Blackburn 1983].) The quasi-realist wants to explain how attitudes that appear committed to a certain metaphysics can be understood in a way that doesn't commit to that metaphysics.⁶

I'll develop the quasi-realist account in stages.

The first stage is to accept 'Cr(θ)' as a *correct* shorthand for ' $\mu(\theta)$ ':

$$Cr(\theta) =_{def} \mu(\theta),$$

Skyrms [1980; Skyrms 1984] calls this a *pragmatic* reduction of chance. Unlike a *semantic* reduction, a pragmatic reduction offers no informative analysis of chance statements. Still, credences about chance are reduced to other aspects of one's belief state. If you are rational and regard a sequence of binary outcomes as exchangeable and extendable, then the measure μ in de Finetti's theorem represents your beliefs about chance.

In stage 2, we have to address a version of the Frege-Geach problem. The measure μ is only defined over the chance parameter. It is not defined, say, for conjunctions of a chance hypothesis and a hypothesis about outcomes. If you're confident that a coin is biased towards heads, however, you may also be confident that *the coin is biased towards heads and there will be a heads among the first 3 outcomes*. How can we make sense of that?

The answer is simple. Remember that de Finetti's representation gives us an implicit credence over chance,

$$Cr(\theta) = \mu(\theta),$$

⁶ In the limit, quasi-realists want to vindicate *all* realist attitudes and judgements, leading to the problem of "creeping minimalism" [Dreier 2004]. I don't think we should go that far. In particular, we should not try to vindicate the realist's implausible mysticism about our epistemic access to the chance facts.

as well as an implicit credence over outcomes conditional on chance,

$$Cr(x_1, \dots, x_n \mid \theta) = \theta^{\#h} (1 - \theta)^{\#t}.$$

We can now define a joint probability over outcomes and chance hypotheses, mirroring the standard rule of probability that $Cr(A \wedge B) = Cr(A) \times Cr(B \mid A)$:

$$Cr(x_1, \dots, x_n, \theta) = Cr(x_1, \dots, x_n \mid \theta) \times Cr(\theta).$$

At this point, it is advisable to introduce a more perspicuous notation. I'll use 'Cr' for the original credence function that is only defined over outcomes, and 'Cr*' for an extended credence function that is defined over outcomes and chance hypotheses. Formally, if Cr is defined over an algebra (Ω, \mathcal{F}) , then Cr* is defined over the product algebra $(\Omega \times \Theta, \mathcal{F} \times \mathcal{B})$, where Θ is the space of chance hypotheses and \mathcal{B} is the Borel algebra on $[0, 1]$. We assume that Cr is exchangeable, and require that for every sequence x_1, \dots, x_n and Borel set $A \subseteq [0, 1]$,

$$Cr^*(x_1, \dots, x_n, \theta \in A) = \int_A \theta^{\#h} (1 - \theta)^{\#t} \mu(d\theta).$$

This means that the extended credence is reducible to the original credence: once you fix Cr (and exchangeability holds), Cr* is fixed as well. But while Cr is only defined over outcomes, Cr* is defined for arbitrary Boolean combinations of outcomes and chance hypotheses.

(The idea might be familiar from research in AI. When reasoning about a given feature space, it can be useful to extend the space with extra parameters and perform reasoning in the extended space. Scientific reasoning takes place in the extended space of Cr*, but it is in principle reducible to reasoning in the original space of Cr.)

This was stage 2. In stage 3, we turn to assertions about chance. So far, we've only talked about credence. (This reflects the Bayesian orientation of de Finetti's gambit. In the Bayesian picture, rational agents are never sure about the chances.)

Since chance statements do not have a reductive analysis, the semantics of such statements must involve an extra parameter: We can't assign a truth-value to 'Ch(H) = .5' relative to a *possible world*, which is just a sequence of outcomes. Instead, chance statements will be true or false relative to a pair of a world and a complete specification of the chances.

Fortunately, we've already introduced that same parameter into the representation of beliefs. The Cr* function is defined over a domain of pairs of outcomes and chance hypotheses. This allows for a simple theory of language use: chance statements can express beliefs, much like statements about outcomes. By saying that a coin is biased towards heads, you might express the attitude $Cr^*(Ch(H) > 1/2) \approx 1$.

To flesh out this idea, we'd need a general theory of assertion and how it relates to credence. This is a big topic that is largely independent of the present project. We might, for example, follow the influential model of [Stalnaker 1978, 2014: ch.2], in which assertions function to update a common ground of propositions that are commonly accepted by the participants of the conversation. Common acceptance means that everyone accepts them, everyone accepts that everyone accepts them, and so on. Accepting a proposition means treating it as certain – acting as if one gave it credence 1 (see

[Dinges 2024]). Traditionally, the propositions in the common ground are modelled as sets of worlds; we would now model them as sets of world-chance pairs. The rest of the story can remain the same.

Under the hood, however, assertions about chance play a very different role from assertions about outcomes. This becomes clear if we look at the un-extended credence Cr of speakers and hearers. If your extended credence is sure that a coin is biased towards heads, your un-extended credence is skewed towards worlds with more heads than tails (although it allows for the possibility of more tails than heads), and this skew is insensitive to evidence about individual outcomes. By conveying that the coin is biased towards heads, you would convey this resilient skew.

What about mixed statements that are partly about outcomes and partly about chance? Their semantics poses no problem. Since we have the extra parameters both in the semantics and in the representation of beliefs, we can handle not only Boolean “mixed” statements, but also attitude reports and modals with occurrences of chance terms.

This was stage 3. I want to mention an optional fourth stage that might help to secure the intuitive objectivity of chance.

To a large extent, our construction of Cr^* already achieves this. Consider, for example, the following statement:

- (1) If we change our beliefs about chance, the chances change.

Is this true according to our extended credence function Cr^* ? No. In the earlier toy construction, Cr^* was only defined over pairs of outcomes and chances, and we didn’t consider beliefs as part of the outcomes. In a more realistic model, the “outcomes” would comprise the entire world – the complete history of all occurrent events. This would presumably include our beliefs. Now we can envisage possibilities in which our beliefs about chance change. Does Cr^* assume that these are possibilities where the chance itself changes? Surely not. This would require a very strange underlying credence Cr . Cr^* may well be sure that the chance never changes at all, and certainly not in response to our beliefs.

Similarly, consider (2):

- (2) If our beliefs about chance were different, the chances would be different.

There is no reason to think that this has non-negligible credence in Cr^* . On the contrary, a plausible rule for how we interpret counterfactuals about ordinary events is that we hold fixed the actual chances. This semantics can easily be adopted in the present framework. It renders (2) straightforwardly false.

In sum, the extended credence function Cr^* already rejects statements expressing a dependence of chance on our beliefs. It similarly allows for the possibility that we can be wrong about the chances, and that we may never know the true half-life of radium-227.

Still, one might feel that there is something disquietingly subjectivist or relativist about de Finetti’s gambit. Imagine two scientists with radically different symmetries in their beliefs about outcomes. As a consequence, they disagree about the chances, even after observing the same outcomes. We can imagine that each of them is writing a textbook. According to scientist *A*, the half-life of radium-227 is 42 seconds. According to scientist *B*, it is 42 *years*. Each side regards the other as wrong. We, too, may take sides, reflecting the symmetries in our own beliefs. But isn’t scientist *B* *objectively* wrong?

It's hard to *say* that de Finetti's gambit makes the chances depend on our whims, but one can't shake the feeling that it does.

Bayesians have a standard response to worries about objectivity. Various "convergence theorems" show that different priors tend to converge when faced with the same data. You'd have to start with *very* wacky priors to end up with an extended posterior credence in which the half-life of radium-227 is around 42 years. The fourth stage of the quasi-realist program is, in essence, to declare such priors irrational.

More precisely, I want to suggest a version of the de Finetti gambit in which we switch from a radical subjectivist interpretation of scientific reasoning (which de Finetti himself seemed to favour) to a more objectivist interpretation. On the objectivist interpretation, there is a correct prior (or a set of correct priors), selected by objective criteria. We might think of this prior Cr_o as a kind of idealization of the scientist's actual prior credence Cr . The equality $Cr = Cr_o$ is a rational norm.

De Finetti's theorem still applies to Cr_o . It allows us to define an extended probability Cr_o^* over outcomes and chance hypotheses. Let's assume that the scientist's actual credence Cr is also defined over the extended space of outcomes and chance hypotheses. (Why? Because people plainly do have beliefs about chance.) The rational norm then becomes $Cr = Cr_o^*$.

On the revised account, we no longer reduce actual credence about chance to credence about outcomes. We don't even assume that actual credence is probabilistically coherent. And we can allow for confused agents whose credences about chance float freely from their credences about outcomes.

Bayes' Rule and the Principal Principle now become norms. But they are not the mysterious norms that they are on a fundamentalist account of chance. They don't dictate how information about one aspect of reality should affect opinions about a metaphysically separate aspect of reality. Rather, they are a consequence of the basic norm $Cr = Cr_o^*$ together with the fact that credences about chance are reducible for *rational* credence functions, in a way that guarantees Bayes' Rule and the Principal Principle.

5 Bounded rationality

I have assumed that (rational) credence functions over outcomes are exchangeable: invariant under permutation. This specific assumption is not required for de Finetti's gambit. De Finetti's Theorem can be generalized to credence functions that are merely "partially exchangeable", and further to credence functions that are invariant under an abstract class of outcome transformations. (See [Diaconis and Skyrms 2017: ch.7] for a summary.) But we seem to need *some* substantive symmetry in the credence function over sequences to get the program off the ground. What justifies this assumption?

Objectivists might say that it follows from the requirement that a rational prior is indifferent between "similar" cases: sequences with the same outcome counts are "similar". But is this true?

Suppose the first few outcomes in a sequence are HTHTHTHTHTHTHTHT. Isn't this sequence more likely to continue with H than with T? If so, a rational prior won't be exchangeable.⁷

Intuitively, if we don't have any information about a sequence, we should divide our credence between different hypothesis classes: one class models the outcomes as independent chance events,

⁷ By exchangeability, HTHTHTHTHTHTHTHT and THTHTHTHTHTHTHTHT are equally likely to continue with H.

another as following a deterministic pattern (like HTHTHTHT...), another models them as a Markov chain in which the next outcome probabilistically depends on the previous outcome.

Such a credence function will have very few symmetries. As far as I can tell, none of the generalizations of de Finetti's Theorem would apply. I suppose we might try to find the required symmetries in the credence function conditional on a particular hypothesis class, but it's not clear to me how the conditional symmetries we now require could be cashed out purely in terms of credences over outcomes.⁸

Of course, this is merely a technical problem: the credence functions over outcomes that can be represented in the way I suggested will have *some*, rather complicated, feature in common. The real question is why we should believe that a scientist's actual or idealized credence function should have this feature. *Assuming that there are no objective chances in the world*, why should a rational credence function be decomposable into a mixture of credence functions that appear to be unsure about the chances?

I want to suggest a speculative addition to de Finetti's gambit that might help answer this question.

Suppose a coin is tossed 50 times. There are $2^{50} = 1,125,899,906,842,624$ possible results. A cognitively unlimited agent could store these one by one, in a database with 2^{50} entries, each paired with a (tiny) probability.⁹ For realistic agents, this approach is infeasible. It would not only require an excessive amount of memory, it would also make updates cumbersome and slow: it takes a lot of time to change a quadrillion floating point records.

How else could we store a probability measure over the 2^{50} sequences? Well, we could represent the outcomes as independent chance events. That is, we could store a distribution μ over a parameter θ , and then compute the probability of any sequence as $\int_0^1 \theta^{\#h} (1 - \theta)^{\#t} \mu(d\theta)$. If we choose μ from the family of beta distributions, this only requires storing *two* parameters (quite an improvement over 1.1 quadrillion), and updating becomes almost trivial: after observing an outcome we simply have to add 1 to one of the parameters. Unlike the previous approach, this method neatly generalizes to sequences of arbitrary length.

If an agent uses this approach, their credences will be exchangeable. Exchangeability, I suggest, may or may not be a norm of ideal rationality, but it naturally arises from bounded rationality.

I've argued above that even moderately rational agents should not have exchangeable credences. If a sequence starts with an obvious pattern, one should suspect that the pattern continues. To allow for this, we need a somewhat more complex representation. The agent's credence might be stored as a mixture of the above exchangeable credence and, say, a credence that's uniform over a certain class of deterministic patterns. Generalizing, the agent might divide their credence between different hypothesis classes in the way I suggested above.

This kind of agent will have more subtle, conditional symmetries in their credences over outcomes. We don't need to worry about how these symmetries could be described. Nor do we need to worry about *deriving* a generalized de Finetti-style representation from the symmetries. We already know that the agent's credence is representable as a mixture of chance hypotheses because we've stipu-

⁸ One problem: if we want to identify hypothesis classes with sets of sequences, where does the alternating sequence HTHTHTHTHTHTHT go? If it doesn't go to the "independent chance event" class, the credences conditional on that class won't be exchangeable.

⁹ Because the probabilities must add up to 1, $2^{50}-1$ records would actually be enough.

lated that that's how it is stored! Beautiful as it may be, in this version of de Finetti's gambit, the mathematics of de Finetti's theorem becomes irrelevant.

I'm not suggesting that our brain stores our beliefs in exactly the way I have described.¹⁰ More plausibly, our brain uses some messy approximation to the described format. Judgements about chance still involve an element of idealisation.

De Finetti's original gambit makes a big deal out of a mere mathematical possibility: an exchangeable credence is *representable* as a mixture of chance hypotheses. My revised gambit suggests that this representation is (to some approximation) *cognitively real*. This might explain why robust realism about chance is so attractive. In our internal model of the world, there really are special chance facts that are only probabilistically related to outcomes. The objective reality of chance is an illusion, created by the way our brain stores and updates information.

6 Outlook

De Finetti's gambit is often presented as a number of mathematical theorems together with a revolutionary philosophical comment: "[physical] probability does not exist!". I've tried to explain how the gambit can be developed into a philosophically more substantive, quasi-realist account of physical probability.

Throughout, I have relied on the simplifying assumption that the world is a sequence of coin flips. Dropping this assumption would lead to a much richer picture. The fuller picture would clarify the connection between chance, projectibility, and resiliency, as discussed in [Skyrms 1980: Ch.IA], [Skyrms 1984: ch.3], and [Skyrms 1994]. It would shed light on the obvious affinity between de Finetti's gambit and the "arbitrary function" approach to chance, discussed (for example) in [von Plato 1983], [Strevens 2003], and [Myrvold 2016]. It would, I think, also help to understand causal propensities in Bayesian networks (see [Pearl 2000: ch.7]) But all these points are largely independent of the issues I have discussed.

One last point. If you are familiar with expressivist and quasi-realist accounts in other domains, especially in ethics, you may be surprised by how smoothly the story worked for chance. We could easily solve the Frege-Geach problem (compare e.g. [Woods 2017]). We did not require a revisionist theory of assertion (compare e.g. [Cuneo 2006]). We had no problem with credence about chance (compare e.g. [MacAskill et al. 2020: ch.7]), and we could immediately explain the apparent objectivity of chance (compare e.g. [Gibbard 1990: ch.8]).

Why did the story work so smoothly? Because we started with credence. Metaethical expressivists, by contrast, generally start with language and assertion, and deal with credence only as an afterthought, if at all. A metaethical expressivism that follows the de Finettian path would start with credences about morality. It would explain how our credences over an extended logical space with an extra moral parameter are determined, perhaps from purely descriptive credences together with more desire-like attitudes (as in [Robinson and Steele 2023] perhaps, but there are many alternatives).

¹⁰ It obviously doesn't because it does not assume that the world is a sequence of binary outcomes. That's not the point I want to make. The point is that our brains at best approximate a mixture of chance hypotheses over the more high-dimensional space of real events.

Only then would it turn to the semantics and pragmatics of moral language. If I could give a piece of advice to metaethical expressivism, it would be this: *credence first!*

References

- José M. Bernardo and Adrian F.M. Smith [2009]: *Bayesian Theory*. John Wiley & Sons
- Simon Blackburn [1983]: “Opinions and Chances”. *Proceedings of the British Academy*, 69: 91–112
- [1993]: *Essays in Quasi-Realism*. Oxford: Oxford University Press
- Terence Cuneo [2006]: “Saying What We Mean: An Argument against Expressivism”. In Russ Shafer-Landau (Ed.) *Oxford Studies in Metaethics*, vol 1. Oxford: Clarendon Press, 35–71
- Bruno de Finetti [1970]: *Theory of Probability*. New York: John Wiley & Sons
- Persi Diaconis and David Freedman [1980]: “Finite Exchangeable Sequences”. *The Annals of Probability*, 8(4): 745–764
- Persi Diaconis and Brian Skyrms [2017]: *Ten Great Ideas about Chance*. Princeton University Press
- Alexander Dinges [2024]: “Assertion and Certainty”. *The Philosophical Quarterly*, 74(1): 169–186
- James Dreier [2004]: “Meta-Ethics and the Problem of Creeping Minimalism”. *Philosophical Perspectives*, 18: 23–44
- John Earman [1992]: *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. Cambridge (MA): MIT Press
- Allan Gibbard [1990]: *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Harvard University Press
- Alan Hájek [1997]: ““Mises Redux” – Redux: Fifteen Arguments against Finite Frequentism”. *Erkenntnis*, 45(2-3): 209–227
- [2009]: “Fifteen Arguments against Hypothetical Frequentism”. *Erkenntnis*, 70(2): 211–235
- [2023]: “Interpretations of Probability”. *The Stanford Encyclopedia of Philosophy*
- Colin Howson and Peter Urbach [1993]: *Scientific Reasoning*. La Salle (Ill.): Open Court Press, 2nd edition
- David Hume [1739]: *A Treatise of Human Nature*. London: John Noon
- Richard Jeffrey [2004]: *Subjective Probability: The Real Thing*. Cambridge: Cambridge University Press

-
- Pierre-Simon Laplace [1814]: *Essai Philosophique Sur Les Probabilités*. Paris: Courcier
- David Lewis [1980]: “A Subjectivist’s Guide to Objective Chance”. In Richard C. Jeffrey (Ed.) *Studies in Inductive Logic and Probability*, vol 2. Berkeley: University of California Press, 263–293
- [1994]: “Humean Supervenience Debugged”. *Mind*, 103: 473–490
- William MacAskill, Krister Bykvist and Toby Ord [2020]: *Moral Uncertainty*. Oxford: Oxford University Press
- Wayne C. Myrvold [2016]: “Probabilities in Statistical Mechanics: What Are They?” In Christopher Hitchcock and Alan Hájek (Eds.) *The Oxford Handbook of Probability and Philosophy*, Oxford: Oxford University Press
- Judea Pearl [2000]: *Causality: Models, Reasoning and Inference*. Cambridge: Cambridge University Press
- Pamela Robinson and Katie Steele [2023]: “Moral Uncertainty, Noncognitivism, and the Multi-Objective Story”. *Noûs*, 57(4): 922–941
- 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
- [1994]: “Bayesian Projectibility”. In Douglas Stalker (Ed.) *Grue!: The New Riddle of Induction*, Chicago: Open Court, 241–262
- David Spiegelhalter [2024]: “Why Probability Probably Doesn’t Exist (but It Is Useful to Act like It Does)”. *Nature*, 636(8043): 560–563
- Robert Stalnaker [1978]: “Assertion”. In Peter Cole (Ed.) *Syntax and Semantics*, vol 9. New York: Academic Press, 315–332
- [2014]: *Context*. Oxford: Oxford University Press
- Michael Strevens [2003]: *Bigger than Chaos*. Cambridge (Mass.): Harvard University Press
- Jan von Plato [1983]: “The Method of Arbitrary Functions”. *British Journal for the Philosophy of Science*, 34: 37–47
- Jack Woods [2017]: “The Frege-Geach Problem”. In *The Routledge Handbook of Metaethics*, Routledge