

Credence—and Chance—Without Numbers (and with the Euclidean Property)

Tim Maudlin
NYU and John Bell Institute

Dedicated to the Spirit and Memory of Detlef Dürr

Abstract

Accounts of both rational credence and of objective chance have always confronted difficulties associated with events that are assigned “probability zero” by the usual Kolmogorov probability function used to model the situation. One sort of solution recommends extending the number field used to represent credences and chances to the surreals or hyperreals. But the correct solution—the solution that always respects the Euclidean property—is to eliminate numbers from the fundamental representation of credence and chance altogether in favor of a system of relations. This solution also sheds light on other paradoxes, such as the Banach-Tarski paradox and the St. Petersburg paradox.

The Problem (?)

Suppose that an infinitely sharp dart is thrown at a circular dartboard in such a way that each individual point might be hit: nothing would prevent the dart from landing there. In other words, suppose that a point is randomly chosen from a disk in such a way that every point might be chosen. Let two of these points be labeled p and q . Here are some propositions about credence and chance that every rational person should accept.

First, the proposition that either p or q will be chosen is strictly more credible—more rationally believable—than the proposition that p will be hit. If someone did not believe the former more firmly or to a higher degree than the latter then we would have no idea how to make sense of their state of belief. We can illustrate our bewilderment by appeal to practical rationality. Suppose the person is offered two choices: a ticket that can be redeemed for an item the person wants if p is hit, or two tickets: the one that pays if p is hit and other if q is hit. And suppose the person is assured that their choice will in no way influence or be correlated with which point is chosen¹. Then anyone who fails to positively prefer getting both tickets to the single ticket would be regarded as insane or as not having understood the situation, or as not wanting the item after all. Because it is, as a matter of rational necessity, more credible that either p or q will be hit than that p will be. We will write this rational relation of credences in these two propositions as $\text{Cr}(p \text{ or } q \text{ is chosen}) > \text{Cr}(p \text{ is chosen})$, where “Cr” stands for the credence in or credibility of the argument. If it is clear that the topic of discussion is credences, we may just write “ p or q is chosen” $>$ “ p is chosen”.

Parallel to this is an analogous situation with respect to objective chance. The objective chance of “ p or q is chosen” being true is strictly greater than the objective chance of “ p is chosen” being true. Or, as we will write, $\text{Ch}(p \text{ or } q \text{ is chosen}) > \text{Ch}(p \text{ is chosen})$. The similar structure of these two inequalities is no coincidence. Indeed, it is an instance of the proper formulation of David Lewis’s Principal Principle.

This little observation can be expanded. Let Σ be any set of points in the disk and Σ' any proper subset of Σ . Then we have both $\text{Cr}(\text{Some point in } \Sigma \text{ is chosen}) > \text{Cr}(\text{Some point in } \Sigma' \text{ is chosen})$

¹ That is, we are not in a situation like Newcombe’s problem.

chosen) and $\text{Ch}(\text{Some point in } \Sigma \text{ is chosen}) > \text{Ch}(\text{Some point in } \Sigma' \text{ is chosen})$. These inequalities are analytic claims about rational relative credibility and relative chance. If someone denied either of them we would infer that the person simply did not understand the notion of credibility or of chance, or was joking or pretending or insane.

Demanding that either credence or chance obey these inequalities for all sets and their proper subsets is called the *Euclidean* property. It is in the same conceptual neighborhood as a property called *regularity*, which demands that all propositions regarded as (epistemically) possible be assigned a numeric measure of credibility greater than 0. The conceptual pedigree of the Euclidean principle can be traced back to Euclid's *Elements*. Among Euclid's Five Axioms—the self-evident truths that apply in many fields including but not limited to geometry—the last is: the whole is greater than the part. (Euclid obviously meant the proper part, since no one had adopted the peculiar convention of calling an entire object a part of itself). In the relevant sense, a proper subset is a part of a set, and therefore the credibility and chance of the chosen point lying in the whole set is greater than the credibility or chance of it lying in the part, so long as every point might be chosen or an agent thinks it might be chosen.

I do not think that the Euclidean property is optional for either rational credence or for objective chance. Any rational person, in such a situation, would have to regard “ p or q will be chosen” as strictly more credible or plausible than “ p will be chosen”, and for any acceptable account of objective chance the chance that either p or q will be chosen must be strictly greater than the chance that p is chosen.

How should the structure of credence or of chance be represented? A person can regard one proposition as more credible than another, which we represent as $\text{Cr}(P) > \text{Cr}(Q)$. They can regard one as at least as credible as the other: $\text{Cr}(P) \geq \text{Cr}(Q)$. If a person regards two propositions as *exactly* as credible as each other we represent that as $\text{Cr}(P) = \text{Cr}(Q)$. If the person is unwilling to accept either $\text{Cr}(P) > \text{Cr}(Q)$ or $\text{Cr}(Q) > \text{Cr}(P)$ or $\text{Cr}(P) \geq \text{Cr}(Q)$ or $\text{Cr}(Q) \geq \text{Cr}(P)$ or $\text{Cr}(P) = \text{Cr}(Q)$, then they have no opinion about which is more credible or plausible than which. In such a case, one would be indifferent between a ticket that pays off if P occurs and a ticket that pays off if Q occurs, but not because one regards them as exactly equally credible. We will represent this attitude as $\text{Cr}(P) \succ\prec \text{Cr}(Q)$. Unlike $=$, $\succ\prec$ need not be transitive. For example, suppose I am asked which proposition is more plausible or credible: that in 2030 the New England Patriots win the Super Bowl or that in 2030 Argentina wins the World Cup. Or suppose I am

offered the choice between a ticket that gives a prize if the first happens and a ticket that gives the same prize if the other happens. I would be indifferent: I just don't know enough or have precise enough degrees of comparative credibility in those propositions. I could happily flip a coin to determine which ticket to take. $Cr(\text{Patriots win 2030 Super Bowl}) \succ\prec Cr(\text{Argentina wins 2030 World Cup})$. But now suppose there is a third ticket that pays if Argentina just gets to the 2030 World Cup finals. Given my state of ignorance, I am indifferent between *that* ticket and the one that pays if the Patriots win. But obviously I would prefer the bet that Argentina gets to the finals to the bet that Argentina wins. That is $Cr(\text{Patriots win}) \succ\prec Cr(\text{Argentina wins})$ and $Cr(\text{Patriots win}) \succ\prec Cr(\text{Argentina reaches finals})$ but $Cr(\text{Argentina reaches finals}) > Cr(\text{Argentina wins})$. So $\succ\prec$ is structurally different from $=$.

A (slightly idealized) credal state is a set of propositions together with a specification of one of these attitudes for each pair of propositions P and Q:

- 1) P is more credible than Q, represented as $Cr(P) > Cr(Q)$ or $P > Q$
- 2) Q is more credible than P represented as $Cr(Q) > Cr(P)$ or $Q > P$
- 3) P is at least as credible as Q, represented as $Cr(P) \geq Cr(Q)$ or $P \geq Q$
- 4) Q is at least as credible as P, represented as $Cr(Q) \geq Cr(P)$ or $Q \geq P$
- 5) P is exactly as credible as Q, $Cr(P) = Cr(Q)$ or $P \approx Q^2$
- 6) I don't have any opinion about which is more credible, represented as $Cr(P) \succ\prec Cr(Q)$ or $P \succ\prec Q$

(I have given two alternative representations of the same state of belief: one referencing the perceived credibility of a proposition and the other just the proposition itself. Since it is simpler to write, I usually use the latter symbolization, except where it is important to distinguish perceived credibility from objective chance.) A *Generalized Credal State* is a set of propositions together with a choice of one of these six for each pair³.

² The symbol \approx is recruited for this relation between the credibility of the propositions rather than $=$ because we do not want to suggest that the propositions themselves are identical.

³ Idea that chance or probability or credence should be ultimately reduced to a set of pairwise comparisons of relative credibility or likelihood is sometimes called a *qualitative* approach, as contrasted with a quantitative approach which makes foundational use of functions from propositions to a number field. Qualitative approaches have a long history, including works by de Finetti, Keynes and Koopman. For a similar approach and more historical references see DiBella 2018. I have not included explicit comparisons with earlier approaches because I was unaware of them, and the ideas in this paper are the working out of a single line of thought. Comparisons and contrasts would interrupt the flow of the exposition.

(In mathematical usage, of course, both $a > b$ and $a = b$ entail $a \geq b$, but we are using these symbols such that at most one of the three expresses the person's state of mind concerning the relative credibility of P and Q. If one thinks that P is precisely as credible as Q then one would normally not express that by saying P is at least as credible as Q: that would be misleading. Similarly if one thinks that P is definitely more credible than Q. It is also worth pointing out, as James Hawthorne (2016) does, that given the single primitive relation \geq with its usual mathematical meaning one can define the rest: $A > B =_{df} (A \geq B) \ \& \ \sim(B \geq A)$; $B > A =_{df} \sim(A \geq B) \ \& \ (B \geq A)$; $A = B =_{df} (A \geq B) \ \& \ (B \geq A)$; and $A >< B =_{df} \sim(A \geq B) \ \& \ \sim(B \geq A)$.⁴ As we will see below, these definitions in terms of a single basic relation \geq imply that there are certain rational constraints on credal states. The formal appeal of having just a single fundamental relation is manifest, but still one should not read too much into it. All the truth-functional operators can be reduced to constructions using just the Sheffer stroke, but the idea that the stroke relation is somehow conceptually fundamental and operators such as negation, conjunction and disjunction conceptually derivative is completely implausible.)

Here's a way to represent a Generalized Credal State: consider a graph in which every proposition is represented and the appropriate relation ($>$, \geq , $=$, or $><$) is drawn (facing the correct way for $>$ and \geq) between every pair of propositions. Such a graph might have infinitely many propositions.

What we now seek are constraints on *rational* credal states. Rationality is, of course, a normative notion. People can actually *be* in irrational belief states, but the hope is that once the irrationality is pointed out they will be motivated to remedy it. And even if they are not, they are subject to the criticism that they are....being irrational. If the irrationality is severe enough, we might even conclude that they do not understand the concepts involved.

An obvious rationality constraint: $=$ should be an equivalence relation. It should, for example, obviously be reflexive. If someone claimed to believe that $P > P$ rather than $P = P$ we would rightly conclude that they do not understand the meaning of ' $=$ ' and ' $>$ '. Similarly, $=$ must be transitive. $><$ need not be, as we have seen.

⁴ Warning: Hawthorne uses the symbol \approx in a different way than I do, employing \approx for $><$ and \approx for equality of credence. I find his choice counterintuitive since the standard use of \approx allows that if $A \approx B$ then A might be greater than B or less: the relation is agnostic and does not assert exactly equality. My use of $><$ also indicates such agnosticism. Hawthorne's does not.

The transitivity and symmetry of \approx follow from the transitivity of the fundamental relation \geq mentioned above together with the definition of \approx . Symmetry is trivial from the definition: $(A \geq B) \& (B \geq A)$ is obviously symmetric in A and B. The transitivity of \approx follows from the transitivity of \geq : If $(A \geq B) \& (B \geq A)$ and $(C \geq B) \& (B \geq C)$, then if \geq is transitive we must have $A \geq C$ and $C \geq A$, and hence $A \approx C$. The reflexivity of \approx similarly follows from the reflexivity of \geq .

If we take the graph just mentioned and erase all of the occurrences of \approx and $><$ and \geq we are left with a directed graph indicating relations of greater credibility. The $>$ relation should be transitive if the person is rational. Therefore the directed graph should be acyclic. If there were cycles then one would have to either break the transitivity or have $P > P$ for some P, which is irrational.

The graph should have some nodes from which arrows only emerge. These represent credal absurdities: propositions one regards as certainly false, such as $1 = 0$ or $1 \neq 1$ or $P \& \sim P$ for some unproblematic proposition P. Nothing can be less credible than these. Since they lie at the bottom of the credal hierarchy, they should all be \approx to each other. We will call the credal value of these propositions *Bottom* because they lie at the bottom of the directed graph. Similarly, a proposition like $1 = 1$ or $P \vee \sim P$ should only have arrows going in. All of these are credally equal to each other, and we call their credal value *Top*.

Anything one regards as epistemically possible (possible as far as one knows) must have a credence greater than Bottom. That is what regularity requires. No matter how implausible or far-fetched or unlikely, it is more plausible than $1 \neq 1$. Any proposition that is neither Top nor Bottom is epistemically contingent: as far as one believes, it could turn out either way.

There are other rationality constraints on a Credal State that are associated with the classical truth-functional logical operators. For negation: if $P > Q$ then $\sim Q > \sim P$. For disjunction: If $(\sim P \& Q) > \text{Bottom}$, then $(P \vee Q) > P$ and if $(\sim P \& Q) \approx \text{Bottom}$ then $(P \vee Q) \approx P$. For conjunction: If $(P \& \sim Q) > \text{Bottom}$, then $(P \& Q) < P$ and if $(P \& \sim Q) \approx \text{Bottom}$ then $(P \& Q) \approx P$.

Since we would like to consider infinitary disjunctions, it is convenient to use a generalized truth-functional operator whose argument is a set of propositions. For any set S of propositions, $\vee(S)$ is true iff at least one proposition in the set is true and $\&(S)$ is true iff every proposition in the set is true. Since we are assuming only classical truth values, each is false otherwise.

It is worth our while to reflect for a moment on this rationality constraint for disjunction: If $(\sim P \ \& \ Q) > \text{Bottom}$, then $(P \vee Q) > P$. The antecedent says that one does not regard the holding of Q and the failure of P to be impossible, or in other words $(\sim P \ \& \ Q)$ is an epistemically possible state of affairs. Therefore, $(P \vee Q)$ might turn out to be true even though P is false. But of course $(P \vee Q)$ is true whenever P is true. Therefore, there are strictly more epistemically possible situations that would make $(P \vee Q)$ true than would make P true, hence $(P \vee Q)$ must be strictly more plausible or credible than P .

I don't see how to deny any of the steps of that argument. Anyone who regards $(\sim P \ \& \ Q)$ as possible but $(P \vee Q)$ not to be more credible than P is irrational. This is just an example of the Euclidean principle: the space of (epistemically) possible conditions that would make P true is a proper part of the space that would make $(P \vee Q)$ true, and the whole is greater than the part, so $(P \vee Q)$ is more plausible than P (assuming $(\sim P \ \& \ Q)$ is epistemically possible).

The full Euclidean Principle is nothing but the generalization of this argument to all sets of propositions and their proper subsets, including infinite sets.

Suppose we have two sets of propositions Σ and Σ' . Consider the three propositions $\vee(\Sigma)$, $\vee(\Sigma')$ and $(\vee(\Sigma) \vee \vee(\Sigma'))$. The last is the same as $\vee(\Sigma \cup \Sigma')$. Now suppose $(\sim \vee(\Sigma) \ \& \ \vee(\Sigma')) > \text{Bottom}$, i.e. $\vee(\Sigma' / \Sigma) > \text{Bottom}$. In words: it is epistemically possible that at least one member of Σ' is true even though no member of Σ is. In that case $(\vee(\Sigma) \vee \vee(\Sigma')) > \vee(\Sigma)$, by exactly the same argument given before. If offered a ticket that pays if $\vee(\Sigma)$ is true and a ticket that pays the same if $(\vee(\Sigma) \vee \vee(\Sigma'))$ is true, it would be irrational not to prefer the latter to the former. It would be irrational not to regard $(\vee(\Sigma) \vee \vee(\Sigma'))$ as more plausible or more credible than $\vee(\Sigma)$. And if $(\sim \vee(\Sigma) \ \& \ \vee(\Sigma')) \approx \text{Bottom}$, then the two propositions are exactly equally credible since one is completely certain that neither can be true without the other being true.

Our initial example of the dart throw is just an example of this, with Σ being the set $\{p \text{ is chosen}\}$ and Σ' the set $\{q \text{ is chosen}\}$. And it obviously also follows that the proposition that the chosen point lies in any set Σ is more credible or plausible than the proposition that it lies in any proper subset since Σ' since Σ / Σ' is non-empty and every point could be chosen, so $\vee(\Sigma / \Sigma') > \text{Bottom}$.

In our darts example, for every point on the disk there is the proposition that the dart lands there. Since the dart could—as far as we know—land anywhere, for each point p “The dart lands on p ” $> \text{Bottom}$. And for each set Σ of points and each proper subset Σ' , “The dart lands in Σ ” $>$

“The dart lands in Σ' ”. The set of such propositions is identical to the power set of the points on the board (if we assign “The dart lands on \emptyset ” = “The dart doesn’t land” \asymp Bottom). The relations of credibility among these propositions generated by our rationality principle are isomorphic to the proper subset relations among all the subsets of points on the target. That—according to the Euclidean Principle—is the only rational attitude to take.

This particular rationally mandated credal state has been considered highly problematic by philosophers. But what’s the problem? It is perfectly clear and unproblematic, as far as it goes. We may want to demand more of a rational credal state given a more detailed statement of the situation. For example, suppose we are told not merely that each point may be hit by the dart, but that each point has exactly the same chance of being hit as any other. Then we want to demand not merely that for each pair of points p and q , “The dart lands at p .” $>$ Bottom and “The dart lands at q .” $>$ Bottom, but that “The dart lands at p .” \asymp “The dart lands at q .”. Indeed, that is the only rational attitude to take. But that attitude, supposedly, leads to terrible problems. Let’s rehearse the standard presentation of the problems.

The standard approach to representing credal states goes via a completely different route than we have followed. Instead of focusing of *relations of credibility among propositions*, it starts out by trying to represent a “degree of belief” in a proposition by a numerical *function*. In particular, the cognitive agent is required—presumably threatened or at least browbeaten—to assign a *number* to each proposition, usually a real number between 0 and 1 inclusive. Sometimes these numbers are characterized as what the agent would propose as “fair betting odds” on the truth of the proposition, i.e. odds at which the agent would be willing to take either side of a bet.

Now there are many immediate objections to placing such a demand on any agent. For example, if you asked me at what odds I would willingly accept either side of a bet on the proposition that Argentina will win the World Cup in 2030, the straight answer is “none”. I know next to nothing about soccer, and if I were willing to take either side of a bet on that I’m sure that people more informed on the subject would take me to the cleaners. Why in the world would I put myself in such a position? I do not regard the thing as *impossible* nor as *certain*, so Top $>$ “Argentina wins the World Cup in 2030” $>$ Bottom. But it’s a far way from that opinion to a precise real-valued credence! Allowing me to express my attitude by a *range* of numerical values is, I suppose, somewhat fairer, but still basically ridiculous. I have literally no idea how to fix such a range.

To that extent, the standard approach fails right out of the box by demanding that the cognitive agent be *more* opinionated about things than she or he is.

But what we now must confront is a failure of the standard approach in the other direction: it also requires that the agent in some cases be *less* opinionated than she or he is rationally required to be. In our dart-with-equal-chance example, we are supposed to assign some real number between 0 and 1 to the proposition “The dart will land on p ”. But since it is rationally required that “The dart lands at p ” \asymp “The dart lands at q ”, it would be rationally required to assign *the same* real number to both propositions, and indeed to assign the same real number to the corresponding proposition for *each* point on the target. But it gets worse. Since “The dart lands at p ” \asymp “The dart lands at q ” and since the two propositions are mutually exclusive, the chance that the dart lands at either one or the other must be *twice* the chance that it lands at p , and the rational credibility of the proposition that the dart lands at either one or the other *twice* the credibility that it lands at p . And the chance that it lands on one of N points must be N times the chance it lands on a single point, and also *greater* than the chance that it land on a single point. But now we are really up the creek. For there simply *is* no real number that can satisfy all of the demands of rationality here. If we choose any positive real number—no matter how small—to represent the credibility of “The dart lands at p ”, then there will be a finite set of points \mathcal{E} such that the credibility of “The dart lands in \mathcal{E} .” will be greater than 1, i.e. higher than absolute certainty. But that makes no sense. On the other hand, if the number that represents the credibility of “The dart lands at p ” is 0, then it will be 0 for all individual points and 0 for “The dart lands in \mathcal{E} ” for all finite sets, even proper supersets of p . But it is rationally required that the credibility of “The dart lands in \mathcal{E} ” be *greater* than the credibility of “The dart lands at p ” in such a case. So you can’t assign a positive real and can’t assign 0. But now you have licked the platter clean: you can’t assign *any* real numbers between 0 and 1 to these propositions which will be related to each other in the way that rationality demands our credences in these propositions be related to each other.

At this point some Clever Dick pops up to tell us all about the surreal numbers or the hyperreal numbers or the infinitesimals or some other “number field” with different properties than the reals. But let’s keep our wits about us. *The suggestion that one must make reference to obscure and counterintuitive number fields in order to represent our credence structure in this very simple situation is absurd.* After all, all of us knows immediately what to say about these propositions: if the points are all equally likely to be hit then “The dart lands at p ” \asymp “The dart lands at q ” for all

pairs of points, and “The dart lands in Ξ ” $>$ “The dart lands in Σ ” whenever $\Sigma \subset \Xi$. The first follows directly from the stipulated equality of chances and the second from the Euclidean Principle. If you can’t *represent* these credence relations by associating numbers with the propositions, then all the worse for trying to do so! Why in the world might one be motivated to do that in the first place? After all, the notion of the plausibility or strength of credence in a proposition has nothing *prima facie* numerical about it.

The Use of Numbers in Representation Theorems

Formal approaches to many topics in philosophy of science and science itself make use of representation theorems. Let’s step back and consider what they are and why they are useful.

Suppose there is a domain of entities one is interested, and the objects in that domain have a structure. Just to be concrete, suppose you run a plumbing business and have a warehouse full of pipes. Each pipe has several important characteristics as far as its use in plumbing is concerned, including its length and diameter. Every pipe has a particular length and diameter, and these quantities furthermore stand in ratio relations to each other. None of that structure, *per se*, is numerical.

There are also some important operations that can be carried out on the pipes: they can be laid end-to-end and soldered to make a long pipe, for example. One can also line up two pipes at one end and then saw off the longer one to match the shorter. We are within our rights to call the first operation “addition” and the second “subtraction”, with the sawn-off piece the “difference” between the two pipes. The idea that addition and subtraction are sorts of operations that can be applied to entities other than numbers is also endorsed in Euclid’s Axioms. Axiom 1 states: “If equals be added to equals, the wholes are equals” and Axiom 2 “If equals be subtracted from equals, the remainders are equal”. The axioms apply to numbers, of course, but also to geometrical objects and, for all practical purposes, to pipes.

Our plumber may need to figure out whether any two pipes in his warehouse can be soldered together to form a pipe of a given length. He can, of course, go to the warehouse and start laying pipes end to end to find out. But that would be an extremely time-consuming and laborious operation, so he seeks to construct a more conveniently manipulable *representation* of the pipes.

The representation will contain precisely the information he needs, but in a form he can access and use more simply.

The first idea that strikes him is to create an *analog model* of the pipes: a collection of small wooden dowels that stand in the very same ratio relations of lengths to each other as the pipes do. Each dowel represents—by stipulation—a particular pipe. All that the representation is supposed to capture are the ratio relations, so the size of the dowels can be scaled down uniformly from the size of the pipes. The arbitrary decision about the length of the first dowel chosen is a *choice of scale* and the freedom to choose any length is a *gauge freedom*. Having chosen the first dowel to represent a particular pipe, though, the length of all the remaining dowels is fixed since the dowels are supposed to stand in the same ratios of lengths to each other as the corresponding pipes do.

This analog system can accurately represent the length ratios of the pipes because it is constructed to do so and because the dowels are capable of standing in exactly the same ratio relations to each other as the pipes do to each other. We here have an accurate analog representation of certain features of the pipes, constructed without a number in sight.

Of course, there is no *a priori* guarantee that the smaller dowels *can* instantiate exactly the same structure as the pipes. If the space is Euclidean then they can because the space is scale-invariant. But if the space were uniformly curved then no model built of the dowels (including, say elbow joints) could precisely represent certain structures built of the pipes. So given the structure of the things to be modelled and the structure of the things doing the modelling one requires a proof—a theorem—showing that the model can serve the purpose required of it. That's the point of a representation theorem: to show that the structure to be represented can be relevantly isomorphic to the structure of the representation that is being used for representational purposes. Representations need not employ numbers, and therefore representation theorems need not be about numbers of any sort.

The dowel-representation of the pipes is, of course, not the most convenient thing to use. Even more convenient would be an abstract, purely mathematical representation. And more particularly—because we know how to perform various operations on them—a *numerical* representation.

And the rest of the story is familiar. If the pipes exist in a Euclidean space, then a convenient mathematical representation of their lengths and ratio structures and of the shapes that

can be built from them is provided by assigning a real number to each pipe representing its length.⁵ There is a gauge degree of freedom in such a representation, as there is in the dowel analog representation, which is resolved by the arbitrary choice of a length to be assigned the number 1: the “unit of length”.

Manipulations and calculations can be conveniently done on the abstract representation. This makes the plumber’s life considerably easier. But aside from the practical advantages, nothing of deep metaphysical import has occurred. The pipes are what they are and have the structure they do and can be put together to form the shapes that they can quite independently of the existence of any representation of them at all. All the representations do is allow us to more conveniently think and communicate about them.

Our present subject matter is not mathematics, it is human cognitive states. People believe things in different “degrees” or “strengths”. How much structure do these “strengths” have? So far, we have only assumed that they—at least sometimes!—are comparable in strength. Some can be stronger than others. Some can be exactly equal to others. Some can also be incomparable to others. These comparative degrees of strength *are* what constitute a credal state on the approach being suggested.

Can we create an abstract mathematical representation of a credal state? Of course we can: we already have done it. We represented a credal state by a graph. And we used that graph to specify various rationality constraints on credal states, such as that the directed sub-graph formed from the $>$ relation must be acyclic. Do we need a representation theorem to prove that the representation is isomorphic in the right way to the credal state? Well, the theorem is one line, since it follows directly from the postulate about what credal states are that the representation captures their structure completely. So we have the subject matter, the abstract representation, and a representation theorem, all without a single number in sight.

How do numbers get in the game? Numbers are convenient because we know how to do all sorts of calculations on them, so there is a natural sort of psychological pressure to use them as representations. In the case of the relative strength of beliefs, we are tempted to try to do exactly what the plumber does: assign a numerical value to each proposition $f(P)$ such that $P \approx Q$ iff $f(P) = f(Q)$ and $P > Q$ iff $f(P) > f(Q)$.

⁵ This representation is, of course, an idealization. Pipes no more have absolutely precise lengths than beliefs have absolutely precise degrees of credence. Once one gets to atomic scale, there is no exact point where the pipe ends.

But what we have seen, in various ways, is that that scheme is simply impossible. First, every number is either equal to, greater than, or less than every other, but some beliefs are incomparable in credence. That particular problem might be overcome by using intervals of numbers rather than numbers as the representation. But we have also seen that the Euclidean Principle cannot be satisfied in such a scheme, at least if it uses the real numbers. The response to that is to gesture at using hyperreals or surreals, or intervals of them. But it is time to stop the madness: we already have a perfectly acceptable abstract representation of a credal state, namely the graph mentioned above. Why the relentless obsession to use numbers instead?

Numbers have certain convenient formal properties. Any pair of real numbers can be added and subtracted and multiplied, and can be divided so long as the divisor is not zero. Again apart from zero, all pairs of real numbers stand in ratio relations. This allows for the formulation of equations using numbers that can be guaranteed (sometimes under a condition, such as no number is zero) to have a solution or solutions. These features make it quite reasonable to seek numerical representations: in terms of such representations, for example, physical laws can be written as differential equations.

But despite all these advantages and attractions, if numbers are simply not available—if they can't do the representational job required—then they can't. C'est la vie. Take the hit and move on.

One way of moving on is to define operations on the subject matter itself—or on an adequate representation of it—that have the same or similar formal properties as addition and multiplication of numbers. In our case, that means defining an addition-like and multiplication-like operation on credences. We already have the wherewithal to do the first and can easily also achieve the second.

“Adding” and “Multiplying” Credences

Addition of numbers is a binary operation. It has the formal features of commutativity and associativity, which mean effectively that there is a unique sum of any finite set of numbers. There is an additive identity: zero. Because equal numbers are actually the same number (equality implies

identity, unlike pipes of equal length), they automatically satisfy Euclid’s requirement that equals added to equals are equal and equals subtracted from equals are equal.

Is there any formally similar binary function on credences? Yes, although the range over which the function is defined is limited. Consider any pair of propositions P and Q. Each of P and Q has a credence which—for us—just means that it occupies a definite position in the map of a credal state. The disjunction $(P \vee Q)$ is another proposition in the same credal state. If the credal state is rational, then there are constraints on the relative credibility of P and $(P \vee Q)$: $(P \vee Q) \geq P$. We have already remarked that. To that extent, disjunction is formally similar to addition of non-negative numbers: for any non-negative numbers p and q , $p + q \geq p$.

In the case of numbers, the equality $p + q = p$ holds only when q is zero. There is no such constraint on credences. In particular, $(P \vee P) \asymp P$ for all propositions P. So if we are going to define an addition-like operation on credences using disjunction, we must have a constraint. But we already know the relevant constraint: P and Q should be mutually exclusive, i.e. $(P \& Q) \asymp \text{Bottom}$. So long as this condition holds, the credence of $(P \vee Q)$ can be regarded as the “sum” of the credence of P and the credence of Q. We can express this with a convenient notational device: $\text{Cr}(P) \oplus \text{Cr}(Q) =_{\text{df}} \text{Cr}(P \vee Q)$ whenever $\text{Cr}(P \& Q) = \text{Bottom}$. Or, dropping “Cr” as is our wont, $P \oplus Q =_{\text{df}} (P \vee Q)$ provided that $(P \& Q) \asymp \text{Bottom}$.

The \oplus operation is evidently commutative and associative because disjunction is and the side condition in terms of conjunction is. Any proposition B with the credal value Bottom is an “additive” identity: for all P, $P \oplus B \asymp P$. What about Euclid’s dictum that equals added to equals are equal? It does obtain *whenever the relevant “additions” are defined*. That is, if $P \asymp R$ and $Q \asymp S$, then $P \oplus Q \asymp R \oplus S$ supposing both are defined. This too is an obvious rationality constraint on credal states: if you regard P as precisely as credible as R and Q as precisely as credible as S, and if you regard P & Q as not possible and R & S as not possible, then should regard the two disjunctions as precisely as credible as each other. (Note: the relation here is \asymp , not $><$. There is no similar requirement for $><$, even though in a sense one is “credally indifferent” between P and Q when $P >< Q$. That sort of indifference is not a judgment of exact equality.)

\oplus also fails to be like addition of numbers in some respects. For any two positive real numbers $p > q$, it is possible to add q to itself some finite number of times and reach a number greater than p . That is precisely the Archimedean property of real numbers (a property the hyperreals fail to have). That property cannot possibly exist for “addition of credences” because

there is a maximal possible credence value: Top. No number of “additions” of credence can take you above Top because there is nothing above Top.

\oplus is a binary operation on propositions, so by iteration it defines a ternary operation $((P \oplus Q) \oplus R)$, quaternary, and so on for all integers. Since \oplus is associative and commutative, these operations are defined for finite sets of propositions (assuming, of course, the constraints for the operation are met). We have already introduced an unrestricted disjunction operation over all sets of propositions, finite and infinite, and we can use it to define a similar unrestricted \oplus operation whose argument is any set of propositions. The statement of the constraint for the \oplus operation to be defined, the analog of the constraint that $(P \& Q) \asymp \text{Bottom}$, is also simple: we demand pairwise mutual inconsistency. For all distinct P and $Q \in \Sigma$, $(P \& Q) \asymp \text{Bottom}$

This definition allows us to “add” infinitely many propositions in our original darts example: let Σ be the set of propositions each of which stating that the dart will land on a particular point in the dartboard. Our original definition of a binary operation falls out as a special case since $(P \& Q) \asymp \text{Bottom}$.

It is worthy of note how easy it was to extend the original binary \oplus operation on propositions to cover infinite sets of propositions. Doing the same trick with real numbers requires definitions of limits of infinite sequences, which brings along many complications. In our case, the extension is simple and straightforward, and never requires us to order or enumerate the elements of Σ . Intuitively, addition is an operation on a set of summands, not an ordered set. Because of the commutativity and associativity, the order shouldn’t matter. In our definition, this is manifestly the case. Furthermore, uncountable sets cannot even be enumerated into a sequence, so taking limits of sequences there is automatically hopeless. Our operation, in contrast, is indifferent about the cardinality of the set. And when people form firm judgements about relative rational credences in the dartboard case, they are also completely indifferent about the cardinality of the set of points in the dartboard. The Euclidean Principle is also completely indifferent about cardinality. So there is strong reason to believe that what people are doing when they make judgments of relative credibility in such a case is more akin to using \oplus than to using $+$.

We now have an addition-like operation defined directly on credences. What about multiplication?

Just as the natural place to look for an analog to addition is disjunction, the natural place to look for an analog to multiplication is conjunction. For any pair of real numbers p and q between

0 and 1, $p \cdot q \leq p$. And for any pair of propositions P and Q, $(P \& Q) \leq P$ (if one is rational). A conjunction is never more rationally credible than a conjunct, nor rationally incomparable to a conjunct. Just as disjunction with another proposition never drives credibility down, conjunction never drives it up.

But just as in the case of addition, we have also a disanalogy: multiplication is defined for every pair of numbers, but the analog is only defined under a side constraint. For any positive number p less than 1, $p \cdot p < p$. But for any proposition at all, $(P \& P) \asymp P$. In the case of addition/disjunction, the constraint could be formulated using concepts already in play. But for the case of multiplication/conjunction we need a new bit of machinery.

Credences of propositions “multiply” under conjunction when the propositions are regarded as being about independent subjects. For example, when sequentially flipping a fair coin a natural assumption is that the flips are “statistically independent”, meaning that the credibility that one flip comes heads is unchanged by information about whether any other flip came heads. Let’s consider this case in more detail.

We already have the resources to express the idea that a coin and a particular flipping process are “fair”: $\text{Cr}(\text{Coin lands heads}) = \text{Cr}(\text{Coin lands tails})$. If in addition we have $\text{Cr}(\text{Coin fails to land head or tails}) = \text{Bottom}$, so one is subjectively certain the coin will show heads or tails, then we can even in an obvious way say that $\text{Cr}(\text{Coin lands heads}) = \text{Cr}(\text{Coin lands tails}) = \frac{1}{2}$. The numerical value $\frac{1}{2}$ is indicated given $\text{Cr}(\text{Coin lands heads}) = \text{Cr}(\text{Coin lands tails})$ and $\text{Cr}(\text{Coin lands heads}) \oplus \text{Cr}(\text{Coin lands tails}) = \text{Top}$. So if we numerically represent Top by the number 1 and Bottom by the number 0 (which we should since Bottom is the identity of the operation \oplus), then we should numerically represent the credence of heads and of tails in the fair coin case by $\frac{1}{2}$.

Now: suppose we flip the coin twice. If we assume that the flips are independent, then the credibility of it landing heads twice is clearly lower than the credibility of it landing heads once. Using subscripts to denote the two tosses, $\text{Cr}(H_1 \& H_2) < \text{Cr}(H_1)$. But our rational credences have a much more detailed structure than that. As we are wont to say, if the coin is certain to land one way or the other on each flip, and the coin and flipping is fair, then the chance of getting one heads is $\frac{1}{2}$ and the chance of getting two is $\frac{1}{4}$. $\text{Ch}(H_1 \& H_2) = \text{Ch}(H_1) \cdot \text{Ch}(H_2)$. Note carefully that in the last equation we have switched from Cr to Ch, from claims about rational credence to claims about chance. We are also assuming that chances are properly represented by numbers. The point is that this is a familiar and unobjectionable sort of calculation about chances that uses multiplication.

According to the Principal Principle the chances should determine our credences in this case (assume no “inadmissible information”).

We are already in a position to make good sense of a parallel claim about credences. We can, for example, accept the following criterion for saying $\text{Cr}(P) = 1/4$: If $\text{Cr}(P) = \text{Cr}(Q) = \text{Cr}(R) = \text{Cr}(S)$ and $\text{Cr}(P) \oplus \text{Cr}(Q) \oplus \text{Cr}(R) \oplus \text{Cr}(S) = \text{Top}$, then $\text{Cr}(P) = 1/4$. Using the same sort of criterion, we can similarly provide a sufficient condition for a proposition to have any credence value $1/N$, and then using \oplus again in the obvious way a sufficient condition for a proposition to have credence value N/M for any integers N and M with $N \leq M$. So our “addition-like” operation on the credences of propositions gives us the resources to attribute “rational-numbered” credences to some propositions.

Back to multiplication. The inference from $\text{Cr}(H_1) = 1/2$ and $\text{Cr}(H_2) = 1/2$ to $\text{Cr}(H_1 \& H_2) = 1/4$ is not analytic. Here, for example, is a condition when it fails: you are fully convinced that the coin about to be flipped is not fair but is double-sided, and you think it exactly equally credible that it has two heads and that it has two tails. So $\text{Cr}(H_1) = 1/2$ and $\text{Cr}(H_2) = 1/2$ but $\text{Cr}(H_1 \& H_2) = 1/2$. In this case the condition for credences to behave in a “multiplication-like” way fails. So what exactly is the condition that has been violated?

The key is to consider subjunctive conditionals or conditional credences. Pure subjunctive conditionals are propositions of the form “If P should be the case (or “were to be the case”), then Q would be the case”. The subjunctive conditional is a close cousin of the much more philosophically discussed counterfactual (or contrary-to-fact) conditional, but the counterfactual presupposes the *falsity* of the antecedent while the pure subjunctive conditional is neutral about the truth value of the antecedent. Nelson Goodman’s *Fact, Fiction and Forecast*, the *locus classicus* for discussions of counterfactuals, makes provision for this. Recognition of the pure subjunctive conditional underlies his discussion of what he calls “semifactuals”: subjunctive conditionals whose antecedents are taken to be true. To a first approximation, if you accept the pure subjunctive conditional “If P were the case, Q would be” and also believe that P is false, then you accept the counterfactual “If P had been the case, Q would have been”. And if you accept the pure subjunctive and also believe that P is true, then you accept Q (and therefore the material conditional ($P \supset Q$)). But accepting the material conditional does not in any way similarly commit one to the corresponding pure subjunctive, even when both P and Q are accepted as true. That is, if you accept the subjunctive conditional and accept the truth of the antecedent then you are

committed to the material conditional, but if you accept the material conditional and truth of the antecedent you are not committed to the truth of the subjunctive conditional.

The pure subjunctive conditional bids you to consider the situation in which the antecedent is true—quite irrespective of whether you believe it to be true, or likely, or unlikely, or whatever—and then opine on the status of the consequent under that assumption. Following a relevantly similar convention, we will represent the subjunctive conditional “If Q were to be the case P would be” as $(P|Q)$. For convenience, we can pronounce this “P given Q”. Now it is obvious that if we regard the coin as a fair, two-headed coin being flipped fairly, then $\text{Cr}(H_1|H_2) = \text{Cr}(H_1)$. That is, if the coin is fair and fairly flipped, then the plausibility of it coming heads on the first flip *given it comes heads on the second* is exactly the same as the plausibility that it comes heads on the first.⁶ Finding out—or positing—the outcome of the second flip makes no difference at all to the plausibility of the first. But if one is convinced that the coin is either double-headed or double-tailed, but unsure which, then $\text{Cr}(H_1|H_2) > \text{Cr}(H_1)$. Indeed, in that case $\text{Cr}(H_1|H_2) = \text{Top}$. (We assume here that we are certain the coin is fair, i.e. regard that proposition as Top. Similarly, if we are certain that the coin is biased in some particular way, for example so it has a 25% chance of coming heads, then $\text{Cr}(H_1|H_2) = \text{Cr}(H_1)$. To think otherwise is to commit the gambler’s fallacy.)

So just as $\text{Cr}(P \& Q) = \text{Bottom}$ provides the condition under which disjunction acts like addition for credibilities, $\text{Cr}(P|Q) = \text{Cr}(P)$ is the condition for conjunction to act like multiplication. $\text{Cr}(P|Q) = \text{Cr}(P)$ says that the plausibility of P is completely unaffected irrespective of whether or not Q happens to be the case. Often (but not always!) this is because one simply sees no connection at all between P and Q, the sort of situation that happens when an interlocutor asserts something that strikes one as completely irrelevant to the topic at hand and the response is “What does that have to do with the price of tea in China?”. If P and Q are on what one takes to be evidentially unrelated topics—winner of the World Cup in 2030 and the winner of the Super Bowl in 2030, for example—then one also judges that $\text{Cr}(P|Q) = \text{Cr}(P)$. Note that this equation does not even require

⁶ “ $\text{Cr}(H_1|H_2)$ ” represents the credence one assigns to the proposition H_1 on the stipulated supposition that H_2 holds. It is therefore a two-place function, as it were, and could equally well be represented as “ $\text{Cr}(H_1)|H_2$ ”, with the slash indicating the antecedent of a pure subjunctive conditional. The formal representation here is not really essential: one needs to reflect rather on the meaning of a subjunctive conditional. If one prefers a more uniform formal apparatus, then $\text{Cr}(H_1)$ can be replaced by $\text{Cr}(H_1|1=1)$, for example, or any other trivial and completely manifest truth. But personally I do not see that the clarity of the concept of subjective credence is enhanced in such a formalism. No one in the world could be in a psychological state of not understanding the concept “The credence one assigns to P” but be enlightened by the explication of it as the same as the credence assigned to P should $1 = 1$.

that one have precise “degrees of credence”: I have no idea how plausible it is that Argentina wins the World Cup in 2030, but I believe that whoever wins the Super Bowl has nothing to do with it.

So here is our first stab at an analog to multiplication of credences, parallel to our analog of addition: $\text{Cr}(P) \otimes \text{Cr}(Q) =_{\text{df}} \text{Cr}(P \ \& \ Q)$ provided that $\text{Cr}(P|Q) = \text{Cr}(P)$.

There is a feature of this definition that is unlike the definition of \oplus and is worthy of our attention. Whereas it is immediate that \oplus is symmetric—the condition for $\text{Cr}(P) \oplus \text{Cr}(Q)$ to exist is the same as the condition for $\text{Cr}(Q) \oplus \text{Cr}(P)$ to exist and they necessarily have the same value—the same is not true for \otimes . It is not that $\text{Cr}(P) \otimes \text{Cr}(Q)$ and $\text{Cr}(Q) \otimes \text{Cr}(P)$ could both exist and yet not be equal to each other, for if they both exist they both equal $\text{Cr}(P \ \& \ Q)$ which is symmetric via the symmetry of conjunction. But it still seems possible that $\text{Cr}(P) \otimes \text{Cr}(Q)$ satisfy the condition for being defined and $\text{Cr}(Q) \otimes \text{Cr}(P)$ not satisfy it. That is, we might have $\text{Cr}(P|Q) = \text{Cr}(P)$, but not have $\text{Cr}(Q|P) = \text{Cr}(Q)$. Is this rationally possible?

In the case of the World Series and the World Cup, each is considered irrelevant to the other, so just as $\text{Cr}(\text{Patriots win} | \text{Argentina wins}) = \text{Cr}(\text{Patriots win})$, so also $\text{Cr}(\text{Argentina wins} | \text{Patriots win}) = \text{Cr}(\text{Argentina wins})$. So $\text{Cr}(\text{Argentina wins}) \otimes \text{Cr}(\text{Patriots win})$ and $\text{Cr}(\text{Patriots wins}) \otimes \text{Cr}(\text{Argentina wins})$ are both defined and are equal to each other. Is there a general argument that for a rational agent $\text{Cr}(P|Q) = \text{Cr}(P)$ iff $\text{Cr}(Q|P) = \text{Cr}(Q)$?

The sticky part about trying to make such an argument is not that you could have $\text{Cr}(P|Q) = \text{Cr}(P)$ while $\text{Cr}(Q|P) > \text{Cr}(Q)$ or have $\text{Cr}(P|Q) = \text{Cr}(P)$ while $\text{Cr}(Q|P) < \text{Cr}(Q)$. At least if neither $\text{Cr}(P)$ nor $\text{Cr}(Q)$ is Top or Bottom, one can argue that would be irrational. The sticky part is trying to rule out the possibility that $\text{Cr}(P|Q) = \text{Cr}(P)$ while $\text{Cr}(Q|P) >< \text{Cr}(Q)$. The possibility of incomparable credences makes for a complicated discussion, and one with no parallels in the usual sorts of approaches.

As far as our definition goes, the simple solution is just to cut the Gordian knot. Our final version of the multiplication-like definition is: $\text{Cr}(P) \otimes \text{Cr}(Q) =_{\text{df}} \text{Cr}(P \ \& \ Q)$ provided that $\text{Cr}(P|Q) = \text{Cr}(P)$ or $\text{Cr}(Q|P) = \text{Cr}(Q)$. Now the defined operator is manifestly commutative just because conjunction is.

\otimes behaves just like multiplication of real numbers between 0 and 1 in all respects, with regard to comparative structure. For example, $(\text{Cr}(P) \otimes \text{Cr}(Q)) \leq \text{Cr}(P)$ and $(\text{Cr}(P) \otimes \text{Cr}(Q)) \leq \text{Cr}(Q)$. Top is the “multiplicative identity”, i.e. $\text{Cr}(P) \otimes \text{Cr}(Q) = \text{Cr}(P)$ whenever $\text{Cr}(Q) = \text{Top}$

because $\text{Cr}(P \& Q) = \text{Cr}(P)$ when $\text{Cr}(Q) = \text{Top}$. Bottom acts like 0 in multiplication: $\text{Cr}(P) \otimes \text{Cr}(Q) = \text{Bottom}$ when either $\text{Cr}(P)$ or $\text{Cr}(Q) = \text{Bottom}$, because in that case $\text{Cr}(P \& Q) = \text{Bottom}$.

Furthermore, having $\text{Cr}(P)$ or $\text{Cr}(Q)$ equal to Top or Bottom is the *only* condition under which $\text{Cr}(P) \otimes \text{Cr}(Q) = \text{Cr}(P)$. Suppose both $\text{Top} > \text{Cr}(P) > \text{Bottom}$ and $\text{Top} > \text{Cr}(Q) > \text{Bottom}$, so the agent regards both P and Q as epistemically contingent: each might or might not obtain. And suppose that either $\text{Cr}(P|Q) = \text{Cr}(P)$ or $\text{Cr}(Q|P) = \text{Cr}(P)$, that is, the obtaining of one does not change the credibility of the other. That means that all of $(P \& Q)$, $(P \& \sim Q)$, $(\sim P \& Q)$ and $(\sim P \& \sim Q)$ are epistemically possible. By construction they are all mutually inconsistent. So $\text{Cr}(P) = \text{Cr}((P \& Q) \vee (P \& \sim Q)) = \text{Cr}(P \& Q) \oplus \text{Cr}(P \& \sim Q) > \text{Cr}(P)$ since $\text{Cr}(P \& \sim Q) > \text{Bottom}$. So just as the product of any pair of real numbers greater than zero and less than 1 is less than either multiplicand, the $\text{Cr}(P) \otimes \text{Cr}(Q)$ is less than $\text{Cr}(P)$ and less than $\text{Cr}(Q)$ so long as both $\text{Cr}(P)$ and $\text{Cr}(Q)$ are greater than Bottom and less than Top.

\oplus and \otimes , whenever they are defined, are structurally analogical to addition and multiplication of reals between 0 and 1 (inclusive) with respect to how comparative values behave, with Bottom being the analog of 0 and Top the analog of 1. These analogies include the following principles (in the simpler notation dropping “Cr”):

If $P > R$ and $Q \geq S$ then $P \oplus Q > R \oplus S$.

If $P > R$ and $Q \geq S$ and Q is not Bottom, then $P \otimes Q > R \otimes S$. If $Q = \text{Bottom}$ then $P \otimes Q = R \otimes S = \text{Bottom}$.

If $P \oplus Q = P$, then $Q = \text{Bottom}$.

If $P \otimes Q = P$, then $Q = \text{Top}$ or $P = \text{Bottom}$.

Whenever the \oplus or \otimes operation is defined, these analogs to addition and multiplication of numbers hold. The existence of this analogous structure explains the overpowering temptation to represent credences with real numbers between 0 and 1, using addition to represent disjunction and multiplication conjunction (under the appropriate constraints, such as $\text{Cred}(P|Q) = \text{Cred}(P)$, where Cred is a function mapping propositions to real numbers between 0 and 1). But as powerful as this temptation is, it must be resisted. Because \oplus and \otimes *do not* behave like addition and multiplication of real numbers when applied to infinite sets.

We have seen how to extend the definition of \oplus to cover the “addition” of arbitrary sets of propositions of any cardinality. The extension of \otimes is formally similar:

$\otimes (\Sigma) =_{df} \&(\Sigma)$ provided that for any two disjoint non-empty subsets Σ' and Σ'' of Σ , either $(\&(\Sigma')|\&(\Sigma'')) = \&(\Sigma')$ or $(\&(\Sigma'')|\&(\Sigma')) = \&(\Sigma'')$.

This definition extends \otimes to arbitrary sets, under the requirement that conditioning on any subset be deemed irrelevant to the credibility of any disjoint subset. This definition yields our original binary operation as a special case since the only disjoint non-empty subsets of $\{P, Q\}$ are $\{P\}$ and $\{Q\}$. We cannot follow the simpler shift of demanding of all pairs of propositions $P, Q \in \Sigma$ that either $(P|Q) = P$ or $(Q|P) = Q$ because one can have a triple of propositions that are pairwise conditionally independent of each other but the third is not independent of the conjunction of the other two.⁷

The condition required for $\otimes (\Sigma)$ to be defined for an infinite set is quite strong. It would hold for an infinite set of propositions all on completely unrelated topics, of course, but specifying such an example would be impossible. Nonetheless, we can formulate an intuitive case if we go back to our dartboard example.

Again, imagine a circular dartboard with a dart thrown in such a way that it is (regarded as) equally likely to hit any point, or such that the credibility of the proposition that it hit any point is the same as for any other point, and all are greater than Bottom. Now: draw the vertical diagonal through the dartboard partitioning it into two pieces. By the perfect geometrical symmetry, we at least plausibly have $\text{Cr}(\text{dart lands in right half}) = \text{Cr}(\text{dart lands in left half})$ (if the dart lands on the diagonal both are false). Now do the same thing, dividing the dart board along the horizontal diagonal. Now we have $\text{Cr}(\text{dart lands in right half}) = \text{Cr}(\text{dart lands in left half}) = \text{Cr}(\text{dart lands in top half}) = \text{Cr}(\text{dart lands in bottom half})$. But also, by symmetry, $\text{Cr}(\text{dart lands in right half}|\text{dart lands in top half}) = \text{Cr}(\text{dart lands in right half})$: that is, whether the dart lands in the top or bottom has no bearing on the credibility of the claim that it lands in the top. So $\text{Cr}(\text{Top half}) \otimes \text{Cr}(\text{Right half})$ is well-defined, and equal to $\text{Cr}(\text{Top half} \& \text{Right half})$.

⁷ Here is an example. Alice wants to send a secret yes/no message to Bob in a secure way. She has two messengers, Charlie and Dora. She and Bob agree that the message will be encoded in the parity of digits given to Charlie and Dora: if they carry the same digit, then the message is “yes” and if they carry different digits, the message is “no”. You regard the message as equally likely to be “yes” or “no”. Alice flips a fair coin to decide whether Charlie carries a 0 or a 1, and then gives the appropriate number to Dora. You know all about the protocol.

If you capture Charlie and read his digit, you do not change your credence at all in either what the message is or what number Dora has. Similarly, *mutatis mutandis*, if you capture Dora. And if you discover what the message is, you still have no clue which digit either of them has. So the three propositions are pairwise independent. But if conditional on any two, then credence in the third becomes Top or Bottom.

Once we see this trick, we also see that it can be repeated indefinitely many times. Divide the dartboard into four quadrants using the 45° diagonals. Call the four resulting right isosceles triangles “Top-Triangle”, “Bottom-Triangle”, “Left-Triangle” and “Right-Triangle”. The credibility of the dart landing in the hour-glass shaped union of Top-Triangle and Bottom-Triangle is the same as it landing in the union of Left-Triangle and Right-Triangle. Further, landing in the union of Top-Triangle and Bottom-Triangle is uncorrelated with landing in the top half and uncorrelated with landing in the right half: $\text{Cr}(T-T \text{ or } B-T|\text{Top}) = \text{Cr}(T-T \text{ or } B-T|\text{Right}) = \text{Cr}(T-T \text{ or } B-T)$.⁸ So if our set Σ of propositions were {“The dart lands in the top region”, “The dart lands in the right region”, “The dart lands in Top-Triangle or Bottom-Triangle”} then the condition for $\otimes(\Sigma)$ to be defined would be satisfied and $\text{Cr}(\otimes(\Sigma))$ would be the credibility of the conjunction of the three claims. Next divide the circle into 8 equal slices by dividing each of the four triangles down the middle and take the union of alternating slices. Then 16 equal slices....

There is no end to the set of propositions about where the dart will land that can be generated in this way: they form a denumerable infinity. And by construction, every subset of that infinite collection is probabilistically independent of every other disjoint subset (for each partition we add one of the two propositions generated by the two parts). So the entire infinite set will satisfy the condition for $\otimes(\Sigma)$ to be defined, and the credibility of $\otimes(\Sigma)$ will be the credibility of that infinite conjunction of propositions. And each member of the set will be equally credible as every other member. And every member will have a credibility greater than Bottom. And every conjunction of two distinct members will be less credible than each of the conjuncts.

But now, once again, the real numbers fail us. For if we assign each member of the set the real value 0, then the value of a conjunction of two will equal the value of each conjunct, violating the Euclidean principle. And if we assign each member any positive real value r , then the only thing we could mean by the value of the conjunction of all of them is the limit as N goes to infinity of r^N . But that limit is either 1, if $r = 1$, or 0 if $1 > r \geq 0$. But neither of these “probabilities” is correct. The credibility of the infinite conjunction is greater than Bottom since it is perfectly possible that the dart land in a location that makes them all true. Since the only real number that can be assigned to Bottom is 0, that means that the credibility of the infinite conjunction must be

⁸ There might be a quibble about the dart landing on a point along the boundary between two regions. In order to remove this quibble, we can suppose that all of the boundary regions are eliminated from consideration in all the propositions: e.g. the relevant proposition is “The Dart lands in the top half but not on any boundary region considered in this set of propositions”, e.g. not on the horizontal Right/Left boundary”. That would settle the quibble.

greater than 0. But no matter how small it is, every positive real number is too big. The credibility of $\&(\Sigma)$ must be less than it.

Once more, the fan of hyperreals or surreals or whatever will pipe up: “I have another number field where that is possible”. But first, no number field will be rich enough to cover all possible cases. But much more importantly, we recur to our main question: why insist on using a number field in a representation at all? We have already specified the relative credibilities of all these propositions—including the infinite conjunctions and disjunctions, no matter the cardinality of the set—in a perfectly adequate way. We even can assign relative credence relations to different infinite sets. For example, let Σ be the infinite set specified above, and Σ' be the same set minus the proposition “The dart lands in the right half”. $\otimes(\Sigma)$ and $\otimes(\Sigma')$ are both well-defined and $\otimes(\Sigma) < \otimes(\Sigma')$. Why? Because $\otimes(\Sigma) = (\otimes(\Sigma') \& \text{“The dart lands in the right half”})$, and $\text{Cr}((\otimes(\Sigma') \& \sim\text{“The dart lands in the right half”}) > \text{Bottom}$. So the credibility of the conjunction is less than the credibility of the first conjunct.

The only plausible ground for *demanding* a representation of degrees of credibility via a number field is that there are certain algebraic operations defined over the number field one wants to make use of. But we have shown how to define operations directly on the relative credibility structure that has all the important properties of addition and multiplication without any of the drawbacks: there is a natural extension of each to sets of any cardinality. The wisest course is to let the numbers go altogether, and intervals of numbers or even fancier numerical constructions with them. Produce the entire theory of degrees of credibility—from beginning to end—without mention of any number fields at all, whether real or otherwise. There is nothing to be lost except a bunch of technical impossibilities and headaches. And there is the Euclidean Principle to be gained.

Ratios

Abandoning numbers does not mean abandoning ratios. The theory of ratios is the topic of Book Five of Euclid’s *Elements*, and much of what Euclid says there is right on target. For example, Definition 3: “A *ratio* is a sort of relation in respect of size between two magnitudes of the same kind”. Credences, of course, are magnitudes (they can be compared as greater and lesser)

of the same kind, as are numbers, lengths of lines, areas of closed figures, etc. And then Definition 4: “Magnitudes are said to *have a ratio* to one another which can, when multiplied, exceed one another”. “Multiplication”, as Euclid has in mind, is just the successive addition of elements of the same magnitude. In other words, two magnitudes have a ratio when a finite number of successive additions of magnitudes equal to the smaller yield a sum greater than the larger. If *all* of the pairs magnitudes in a collection of magnitudes have this property then the collection is called *Archimedean*, and there is a ratio between every pair of magnitudes. The positive integers and positive rationals and positive reals are Archimedean collections of magnitudes, as is the collection of closed line segments in the Euclidean plane. The latter example, of course, demonstrates that ratios can exist perfectly well among magnitudes that are not numbers. If we add infinite lines to the collection of lines then it is no longer Archimedean because no finite number of additions of equal line segments can exceed an infinite line. So there is no difficulty at all having a collection of magnitudes which is not Archimedean—some elements stand in no ratio to others—but still which have perfectly precise ratios among some pairs of elements.

It is common to represent ratios by real numbers. For example, it is commonly said that π , which is the ratio of the circumference of a Euclidean circle to its diameter (the ratio of a length to length) is the real number 3.14159.... But that is just loose talk. The accurate locution is this: the ratio of the circumference of a circle to its diameter is *proportional* to the ratio of the real number 3.14159... to the real number 1. Proportionality of ratios is expressed using a notation of colons. In the case of π the proper statement is: circumference of circle:diameter of circle:: 3.14159... :1. The double colon represents proportionality of the ratios, and on each side of the double colon lies a ratio between a pair of magnitudes of the same sort.

When one indicates a ratio by providing just a single real number, what is meant is the ratio between that number and the unit. It’s a harmless enough ellipsis, but for our purposes it is important to be technically and conceptually precise. So we will be fastidious in referring to proportionality of ratios using the classical notation.

One set of magnitudes can have a richer structure of ratios than another. The most famous discovery of this phenomenon is mistakenly referred to as “the discovery that the square root of 2 is irrational”. The Greeks would have no idea what is being referred to by the “square root of 2”. 2 is a number, which means for the Greeks a positive integer. The “square root” of a number N is defined as a number which, when squared, yields N. So the Greeks would say that there is no

square root of 2. No one could refer to a *number* as “the square root of 2” until the invention or discovery of the full set of real numbers. But that was millenniums later.

So if the great apocryphal discovery of Hippasus of Metapontum was not that “the square root of 2 is irrational”, what was it? It was the discovery that the diagonal of a square stands in a ratio to its side that no integer stands to any other integer. Or in other words, there is no pair of integers M and N such that: diagonal of square:side of square:: $M:N$. The structure of ratios among line segments in the Euclidean plane is inherently richer than the structure of ratios among the integers.

(Not to devolve too far into history here, but it seems quite likely some of the Greeks expressed ratios in *anthypharetic* form, which is related to the more familiar continued fractions. In anthypharetic form a ratio is expressed as a sequence of integers. All ratios of integers correspond to finite sequences. The ratio of the diagonal to the side of a square is expressed by the infinite sequence 1, 2, 2, 2, Obviously, the simplest such ratio that is not a ratio of integers is 1, 1, 1, 1, ... That is also known as *The Golden Ratio*, ϕ . Note that while the improper locutions “ $\sqrt{2} = 1.4142135623\dots$ ” and “ $\phi = 1.61803398875\dots$ ” make $\sqrt{2}$ and ϕ appear rather random and arbitrary, “ $\sqrt{2} = 1, 2, 2, 2, \dots$ ” and “ $\phi = 1, 1, 1, 1, \dots$ ” have precisely the opposite effect.)⁹

What we now know is that the structure of rational credences over a set of propositions can be non-Archimedean. For example, there is no ratio between the rational credence in the proposition that the dart lands on a particular point p and the proposition that it lands on the right side of the target. But the rational credence in each of these propositions can stand in perfectly good ratios to that of other propositions. For any other point q , for example, $\text{Cr}(\text{Dart lands on } p \text{ or } q):\text{Cr}(\text{Dart lands on } p):: 2:1$. And plausibly $\text{Cr}(\text{Dart lands in top half}):\text{Cr}(\text{Dart lands in right half}):: 1:1$. (The reason for the slightly hesitant “plausibly” will become clear anon.) Even more plausibly, for a rational person $\text{Cr}(\text{Dart lands in top half}):\text{Cr}(\text{Dart lands in right half}):\text{Area of top half}:\text{Area of right half}$. And most plausibly of all, as an instance of the Principal Principle, $\text{Cr}(\text{Dart lands in top half}):\text{Cr}(\text{Dart lands in right half}):\text{Ch}(\text{Dart lands in top half}):\text{Ch}(\text{Dart lands in right half})$, that is, the ratio of the rational credences is proportional to the ratio of the chances of the events.

⁹ A wonderful account of all this is provided by D. H. Fowler in *The Mathematics of Plato's Academy* (1987). A shorter presentation appears in my *New Foundations for Physical Geometry: The Theory of Linear Structures* (2014), Chapter 7.

The reader may be forgiven for not seeing any daylight at all between these last three claims, but at least the difference in metaphysics is clear. The first asserts that a ratio between credences is proportional to a ratio between numbers, the second that it is proportional to a ratio between areas and the last that it is proportional to a ratio between chances.

In order that ratios be defined between at least some pairs of elements of a field, all that needs to be utilized is a notion of equality and a notion of addition. As far as credences go, these are provided on the field of propositions by \simeq and \oplus . Whether the structure of ratios so defined allows them to be proportional to ratios in some other field—a field of numbers or of lengths or of areas or of chances—is a question that must in each case be investigated in detail. In carrying out such an investigation, we must always keep the cautionary example of Hippasus in mind. He discovered that the universe of ratios among lengths of lines in Euclidean geometry is richer than the universe of ratios among integers. Or as we would say, the universe of ratios among the integers is impoverished compared to the universe of ratios among the real numbers. According to myth, the discovery came as such a shock to the Pythagoreans that he paid for it with his life. Hopefully the price one might pay for such a suggestion is no longer quite so drastic. I have already insisted that the structure of ratios among rational credences cannot be modelled by the structure of ratios among any reals. That is hardly a novel proposal. But we still have to investigate ratios between areas (and in general what areas *are*) and ratios between chances (and what chances *are*), and we should at least be prepared for some surprises.

Bayes' Theorem

One of the keystones—if not the keystone—to the theory of rational credences has long been taken to be Bayes' Theorem. At first blush, this ought to appear quite strange since the theorem itself (unlike, say, the Pythagorean theorem) is a complete triviality, a literal one-liner. Things are not quite so simple, but let's work up to the complications.

Bayes is working in a setting which presupposes a *credence function* in the usual sense: a function f that maps propositions to the real numbers between zero and one. Further, he presupposes that for every pair of propositions P and Q in the domain of the function there is a

conditional $(P|Q)$ in the domain of the function, or at least there is if $f(Q) > 0$. And he requires that $f(P \& Q) = f(P|Q) \cdot f(Q)$ whenever all three terms are defined. That's it.

From these propositions it immediately follows that $f(Q|P) \cdot f(P) = f(P \& Q) = f(P|Q) \cdot f(Q)$. Dropping out the middle term leaves $f(Q|P) \cdot f(P) = f(P|Q) \cdot f(Q)$. That's Bayes' Theorem.

Since the values of the function f are real numbers, if $f(P) > 0$ we immediately have $f(Q|P) = \frac{f(P|Q) \cdot f(Q)}{f(P)}$. This is a trivial consequence of the presuppositions in Bayes' approach. So what's the big deal?

The deal big becomes apparent when we let Q be a speculative *hypothesis* H and P be some uncontroversial *evidence* E . (It is called "evidence", of course, because it is evident, i.e. uncontroversial at least to those with direct empirical access to it.) For example, having lost quite a lot of money playing craps with Shady Sam, one might form the hypothesis that the dice he is using are not fair: they are loaded or shaved or somehow or other monkeyed-with so the chances of the various sides coming up are not the same. That's a vague hypothesis in some respects because the exact method of cheating is not even specified. Therefore providing convincing evidence concerning it might seem to be a hopeless task. You can check for shaved edges, for example, but not finding them does nothing to rule out the dice being loaded or gimmicked in a way one has not even thought of.

Although the hypothesis is vague in one respect it is perfectly precise in another: it asserts that the chances of the various sides are not the same. Shady Sam, of course, insists the opposite: the chances are the same, or at least any difference (due to small imperfections, say) is so small as not to make any difference over the course of the game. So how do we acquire relevant evidence to decide between Shady Sam's hypothesis and yours?

Let's focus on Shady Sam's, since it is more precise. How do we test the hypothesis that the dice are fair other than by checking for all the methods of making them unfair that we can think of? Obviously, by throwing them many times and recording the outcomes.

Now at first glance that method seems like it cannot be of any use at all: if the dice are in fact fair, then any sequence of outcomes is possible, and indeed any sequence is as likely as any other equally long sequence. So you already know *a priori* that no outcome is *inconsistent* with the hypothesis. A naïve sort of Popperian might say that no sequence can falsify the hypothesis, so this method is of no use in testing the hypothesis. Nonetheless, we know perfectly well that if we throw a die twenty times and it comes up six each time we will rightly conclude that the die is

not fair and that the evidence against Shady Sam is strong. What Bayes' Theorem does is give us a handle on understanding that reasoning.

The essential observation is that while what we really want to evaluate is $\text{Cr}(H|E)$ —the credibility of the hypothesis conditioned on some piece of evidence—it is not immediately obvious how to determine what that credence should be. But on the other hand, the converse conditional— $\text{Cr}(E|H)$ — can be straightforwardly calculated. $\text{Cr}(20 \text{ sixes in a row} | \text{die is fair}) = 1/6^{20} = 1/3,656,158,440,062,976$, a quantity commonly known as a miniscule or negligible chance. If you are offered a prize if you throw a fair die 20 times in a row and it comes up six each time, don't even bother. It ain't gonna happen. (This is an example of Cournot's Principle, as we will see.) And if this were the result of throwing Shady Sam's die 20 times in a row we would rightly (i.e. rationally) conclude that the die is not fair and start preparing some cement galoshes.¹⁰ But still: what exactly is the reasoning that underlies this?

After all, Shady Sam insists, *whatever* sequence happened to come up after 20 throws would have been *equally* unlikely under the hypothesis! So you already knew, before you even threw the die, that the outcome would be tremendously unlikely. Why should the fact that it happened to be *this* particular unlikely outcome doom him to sleep with the fishes?

And now Bayes' Theorem comes to the rescue. Before we throw the die, we are unsure which of these two hypotheses is true H_1 : The die is (effectively) fair or H_2 : The die has been gimmicked in some way to give Shady Sam an advantage playing craps. These are the only live hypotheses for us, the only hypotheses credible enough to even worry about. We start, before the test, with some suspicions but those are just that: suspicions. They mean that the antecedent credibilities of H_1 and H_2 are each above, say, $1/1,000$, so the ratio of their credences is at least that much. Our antecedent credibility for the proposition that the die will come 6 all 20 times is quite small. On the supposition that the die is fair, the likelihood is negligible, but on the supposition it is gimmicked the credibility goes way up. Since there is a non-negligible credibility that it is gimmicked, there is a non-negligible credence that it will come all sixes (even if that particular credibility is quite low, since we don't even know how it is gimmicked, if it is).

¹⁰ Note for historical accuracy: organized crime members never actually used "cement galoshes" or "concrete overshoes" for the purpose of committing murder, for rather obvious logistical reasons (as anyone who has had to wait for concrete to set will understand). A complete discussion of the matter can be found at <http://www.todayifoundout.com/index.php/2018/12/did-mobsters-ever-send-people-to-sleep-with-the-fishes-wearing-concrete-shoes/>.

So initially we have two competing hypotheses: the die is effectively fair and the die is gimmicked in a way to help Shady Sam to win. The ratio of the initial credibilities of these hypotheses is no less than 1:1,000. The conditional credibility of $(E|H_1)$ is $1/3,656,158,440,062,976$. The conditional credibility of $(E|H_2)$ is much, much greater than $1/3,656,158,440,062,976$. For example, suppose the die is gimmicked so that the chance of throwing a six is $1/3$ rather than $1/6$. Then the chance of throwing 20 sixes in a row is 2^{20} (i.e. 1,048,578) times greater. And of course, if the die is gimmicked to always come six then the conditional chance is 1. So let's say that $f(E|H_1):f(E|H_2)::1:1,000,000$. And let's say that that initially $f(H_1):f(H_2)::1,000:1$, so one is initially quite unsuspecting of Shady Sam. Now Bayes' Theorem does its magic. Since $f(H_1|E) = \frac{f(E|H_1) \cdot f(H_1)}{f(E)}$ and $f(H_2|E) = \frac{f(E|H_2) \cdot f(H_2)}{f(E)}$, it follows that $f(H_1|E):f(H_2|E)::\frac{f(E|H_1) \cdot f(H_1)}{f(E)}:\frac{f(E|H_2) \cdot f(H_2)}{f(E)}::f(E|H_1) \cdot f(H_1):f(E|H_2) \cdot f(H_2)::1:1,000$. So even if the credibility of H_1 is initially a thousand times greater than that of H_2 , the credibility of $(H_1|E)$ is a thousand times less than $(H_2|E)$, exactly because E is a million times less likely on H_1 than it is on H_2 .

Everyone knows that throwing the die and having it come up six 20 times in a row is absolutely convincing evidence against Shady Sam. Bayes' Theorem explains why. Bravo.

In the little calculation done above, there is use made of multiplication. The quantities $(E|H_1) \cdot f(H_1)$ and $(E|H_2) \cdot f(H_2)$ appear. But it is easy enough to remove all mention of multiplication from the proceedings. Our original statement of the theorem was $f(Q|P) : f(P) = f(P|Q) : f(Q)$, with f being a function into the reals. This is trivially rearranged into $\frac{f(Q|P)}{f(Q)} = \frac{f(P|Q)}{f(P)}$, assuming $f(P)$ and $f(Q)$ are both greater than zero. And this form, in turn, is easily rewritten as a proportionality of ratios: $f(Q|P):f(Q)::f(P|Q):f(P)$. Since $f(P)$ is supposed to be a representation, using real numbers, of the strength of credence in P , we can further generalize this principle as $Cr(Q|P):Cr(Q)::Cr(P|Q):Cr(P)$. And voilà: we now have Bayes' Theorem in a form that can be applied to the credences directly, rather than to their representations, so long as the ratios between the relevant credences are defined.

Let's see exactly what this form of the theorem says. It says that the ratio of the conditional credence $Cr(Q|P)$ to the unconditional credence $Cr(Q)$ is proportional to the ratio of the conditional credence $Cr(P|Q)$ to the unconditional credence $Cr(P)$. When Q is a hypothesis and P some evidence, it means that the credibility of the hypothesis conditional on the evidence stands in the

same relation to the unconditional credibility of the hypothesis as the credibility of the evidence conditional on the hypothesis stands to the unconditional credibility of the evidence. The tremendous utility of this observation lies in the fact that the credibility of the evidence conditional on the hypothesis can often be *calculated*, as the chance of the die coming up six 20 times in a row can be on the hypothesis that the die is fair. So at least in some cases, that quantity is nailed down. What Bayes' theorem now tells you is that if the credibility of the evidence conditional on the hypothesis is much higher than the unconditional credibility, then the credibility of the hypothesis on the evidence should be equally much higher than the unconditional credibility of the hypothesis. Or, in other words, if a hypothesis makes a *surprising* prediction (the hypothesis renders likely a proposition with low initial credence), then the credibility of the hypothesis conditional on the evidence should be much higher than the unconditional credence.

If we add to this that a rational reaction to acquiring evidence (e.g. to suddenly adjusting one's credence in the evidence so it is very high, if not Top) is to make one's new credence in a proposition one's old credence in the hypothesis *conditional* on the evidence, then we have a Bayesian dynamics of credences. Coming to believe in a piece of evidence one initially gave low credence requires one to boost the credence in every hypothesis that predicts that evidence commensurately, where "commensurately" is cashed out using proportionality of ratios.

The moral of this is that rejecting the use of real numbers—or any numbers at all—as representatives of credences does not entail rejecting Bayes' Theorem or all of the interesting explanations of scientific reasoning that rely on Bayes' Theorem. Although the theorem is standardly presented in terms of a function into the reals $f(P)$ and the operations of multiplication and division, it can just as well be formulated directly in terms of credences themselves—or some other representation of credences—and proportional ratios. Abandoning numbers altogether does not mean abandoning the theorem.

Bayes' Theorem stated in terms of proportions of ratios also has immediate consequences for judgments of relative credibility. Once again, the theorem is $Cr(H|E):Cr(H)::Cr(E|H):Cr(E)$. It follows that if $Cr(E|H) > Cr(H)$, i.e. if supposing H renders E more credibly than it presently is, then $Cr(H|E) > Cr(H)$, i.e. the credibility of H on the supposition of E is greater than its credibility presently is. And similarly, of course, if the hypothesis renders some possible evidence less credible. It is a rationality constraint on all rational credal states to obey this rule. If we further

postulate that updating credences by conditionalization is at least rationally *permitted* (if not rationally *required*) we have one rationally defensible way to respond to new evidence.

Credence and the Problems of Old Evidence and New Theories

Bayes' Theorem has many explanatory successes, at least some of which can be recovered in terms of ratios among credences, as we have just seen. But the standard approach to Bayes' Theorem, using probability functions, also confronts challenges. Some of these challenges evaporate once we operate directly with credences rather than with numerical functions.

Two sides of one and the same problem go under the names "The Problem of Old Evidence" and "The Problem of New Theories". The sort of Bayesian conditionalization we just discussed occurs when there is a proposition—the statement of the evidence—which suffers a sudden and dramatic increase in credibility, usually due to observation. Before throwing the die, our credence in "The die comes 6 twenty times in a row" is quite low, and after throwing it is quite high. The question is what effect that change in credence should have on our other credences. Bayesian conditionalization provides a rule.¹¹

¹¹ Bayesian conditionalization is a rational response to the change in credence, but that does not imply that it is the *only* rational response available, or that credence shifts that are not Bayesian are automatically unacceptable. There is nothing objectionable, for example, about a certain amount of *credal drift*. Suppose that, in discussing the probability of an upcoming election, I am mildly inclined to think it more credible that Candidate T wins than loses, $Cr(T \text{ wins}) > Cr(T \text{ loses})$. I discuss this with a friend who is inclined to have the opposite view: $Cr(T \text{ loses}) > Cr(T \text{ wins})$. Say that I am *pessimistic* and my friend is *optimistic* (because T is a complete disaster). We don't disagree on any of the evidence (polls, etc.), but considering the evidence we come down in different places. Since we don't disagree about any known facts, we "agree to disagree": each acknowledges the other as rational, as having a defensible view, even though it differs from one's own.

Now: suppose the next morning—having gotten no new evidence at all—we wake up with reversed moods: I am now optimistic and my friend pessimistic. My credences have changed, but not by conditionalization on anything. That change cannot be modelled by Bayesian conditionalization. But have I done anything rationally *objectionable*? That does not seem possible. Yesterday, I acknowledged that the optimistic attitude was rationally acceptable even though I did not hold it. Today I hold it. So my present position is rationally acceptable: I have done nothing epistemically wrong.

Bayesian conditionalization is always a possible rational way to update credences, but it is not the only rationally acceptable way, and is not rationally mandatory. Confronted with some new evidence, for example, one could change to new conditional credences instead of conditionalizing using the old conditional credences. This is just an instance of the familiar Duhem/Quine thesis.

It is also notable that dealing with relative credences, together with the usual logical operators, allows us to define a state of *weak belief*. S weakly believes $P =_{df} Cr(P) > Cr(\sim P)$ for S. If I and my friend have only negligible credence

But what if we have to react not to some new *evidence* but to the discovery of a new *theory*, a theory we never had thought of, never had considered, never had formed any credence about *at all*. How should the introduction of that novel proposition into the set of propositions about which we have credences change our present credal state?

On the present approach, the first thing we must do is rank the new proposition with respect to other propositions that we already have credences in. The most immediately important judgments to be made concern propositions on relevant or related topics. An example will help.

Gustav is an early 20th century physicist. He is familiar with both Maxwellian electrodynamics and Newtonian gravitation. He sees problems with both. For Newton, there is the apparent action-at-a-distance and instantaneous nature of the gravitational force. There is also the puzzle of the anomalous advance of the perihelion of Mercury, but Gustav figures that something not so surprising—like an unknown planet—may account for that. Overall, the predictive success of Newton’s theory is extremely impressive.

Gustav has much more trouble with Maxwell’s theory because it seems to have to postulate a lumeniferous ether, but no mechanical model he can think of would endow that model with the right properties. He knows about the negative result of the Michelson-Morley experiment. All in all, Gustav judges that $Cr(\text{Newtonian gravity}) > Cr(\text{Maxwellian electrodynamics})$.

In 1915, Gustav reads about the General Theory of Relativity. The Eddington eclipse observations are in the future, so there is no new data, but the anomalous advance of the perihelion of Mercury drops out correctly from the equations without them having been fine-tuned to that

that the election will fail to have a winner, then in our optimistic moods we weakly believe that T will lose and in our pessimistic moods that T will win.

Weak belief seems like a necessary condition for anything deserving the name “belief”. One might, of course, frame a stronger condition. For example, by adding a new relation \gg for “much greater than”, we can define strong belief in P as $Cr(P) \gg Cr(\sim P)$.

In a three-way election, I might have no weak belief about exactly who will win, but still have definite relative credences. Thus, I might say “I don’t know who will win, but if I had to bet I would pick B”.

Radical skepticism is sometimes characterized as a state of no belief. If what one has in mind is weak belief, that cannot be a rational state of mind. If one has in mind strong belief, then a rational person might have no beliefs except in tautologies or other analytic truths. If one makes a distinction between positive and negative propositions—with negative propositions being the negations of positive ones and vice versa—then one could have no weak positive beliefs save in analytic truths, but still make definite judgments that one positive proposition is more credible than another. One way of understanding the aim of the Pyrrhonian Skeptic is to produce a credal state in which for every pair of propositions P and Q $Cr(P) \gg Cr(Q)$, or as they say “No more P than Q”. By our lights, that could not be rational.

data, and the gravitational influence is no longer instantaneous in that theory. So on consideration, Gustav embeds the General Theory into the set of propositions in his credal state in part like this: $\text{Cr}(\text{General Relativity}) > \text{Cr}(\text{Newtonian gravity})$. But furthermore, having demoted Newton, he now judges $\text{Cr}(\text{Maxwellian electrodynamics}) > \text{Cr}(\text{Newtonian gravity})$ reversing his earlier view, and also $\text{Cr}(\text{General Relativity}) > \text{Cr}(\text{Maxwellian electrodynamics})$. So the new proposition is fit into his credal state, both requiring judgments of relative credibility with respect to propositions already there and, in some cases, altering the credence relations between propositions he had already made judgments about.

That is supposed to be both the problem of new theories and the problem of old evidence. The problem of new theories because one has to assign a credence to the novel proposition, and the problem of old evidence because the changes in credence that depend on the evidence flowing from the advance of the perihelion of Mercury cannot be modelled as due to conditionalization. That's because Gustav *already knew* about the advance and so had *already conditioned* on it. Conditioning a second time would not change anything.

But what should be apparent from our perspective is that there seems to be no problem at all. Gustav has some new ideas to think about and fits them into his credal state, making a series of adjustments of relative plausibility. So what?

Aside from there not being any obvious model of this change by standard Bayesian conditionalization, there is an additional problem in the standard approach. The standard approach represents credences by a *probability measure* over propositions, and that measure is always normalized to a specific value: 1. So if Gustav wants to assign some credence to the new proposition, that credence has to *come from* somewhere. He has to dip into the credence assigned to some other proposition and reduce it to free up enough to give to the new proposition. And the question comes up of just where that stock of credence will be drawn from.

Before addressing this question, we must note that on our approach no parallel problem ever arises. New propositions are fit into the credal state by specifying their relative credences to other propositions, and other adjustments to the credal state may follow, but there just is no “pot of credence” that needs to be tapped. That idea comes from deciding to represent credence by a probability function, which was not a good idea in the first place.

Bayesians can take either of two approaches here. One is to insist on a postulate of *logical omniscience*, in the sense that the ideally rational Bayesian agent never has to accommodate a new

proposition or hypothesis because the agent is aware of all possible hypotheses at all times. There is never call to “fit in” a new proposition. The ideal agent knows about the equations of General Relativity and all of their consequences *ab initio*, and always has assigned them a credence.

Of course, if true then it would be a puzzle how an ideal Bayesian agent could ever assign significant credibility to Newtonian gravity in the first place: all the evidence the agent ever has can be modelled with at least as much accuracy as Newton by General Relativity, and GR is in many ways more elegant and simpler. And if the Bayesian ideal agent, being omniscient, always is aware of all possible theories, then the Bayesian agent just seems not be relevant in many respects to understanding actual humans, who are not omniscient! Humans *do* have to confront novel ideas, and much of the history of science is the history of how they do so. So if the Bayesian has to idealize the rational agent in this way, then there are huge issues in the study of scientific rationality that Bayesianism cannot address.

Another tack that the Bayesian can take is not to presume omniscience but to provide the agent with a sort of slush fund of credence, in a catch-all category called “None of the Above”. That is, having only a finite amount of credence to dole out, the agent reserves some for “Some hypothesis, I know not what”. Then when a novel hypothesis is proposed, the agent can dip into that fund to supply it with credence.

But that raises the question of how much credence should go into the slush fund. Make it too high, and the agent perforce becomes a sort of skeptic, always judging it more credible that the truth is something not yet thought up. That might be a reasonable position to take on some subjects, but having it basically forced on one by the formal apparatus is a mistake.

So by shifting from probability measures to relations of relative credence as representations of cognitive states we completely avoid some sticking points of the standard Bayesian approach. But by rewriting Bayes’ Rule itself in terms of ratios of credences we can keep the main motor of many Bayesian explanations.

A little analogy may be of some use here. Formal philosophical theories of things like human states of belief are clearly idealizations. They postulate different structures than the psychological states actually have. But one has to distinguish two sorts of idealizations.

One sort is a simplification that leaves out detail but is still basically on the mark. For example, when calculating the orbit of the Earth or of the Moon, one might idealize the Earth as a uniformly dense perfect sphere. That is a very simple sort of model to do calculations with, and

the predictions in many circumstances will be quite accurate. But one still expects corrections to have to be made from time to time when the idealization fails.

When highly accurate pendulum clocks made in Amsterdam were first imported to the American colonies, they ran slow. It was thought that they had been damaged in transit, and so were sent back to be fixed, but they showed no retardation in Amsterdam. After a while, the problem became clear. The period of a pendulum clock is determined by the length of the pendulum and the strength of the local gravitational field. Using the idealization that the Earth is a uniformly dense perfect sphere, a clock should tick at the same rate everywhere on the surface.

But in fact, due to its rotation, the Earth is not a sphere: it is an oblate spheroid. In going from Amsterdam to the colonies, the clock changed latitude, and therefore increased distance from the center of the Earth, and therefore got into a weaker gravitational field. So it had to run slow. The idealization, in this case, fails, and the failure can be accounted for. But it is still a very good idealization.

In contrast, there is a joke that goes like this:

So a biologist, engineer and physicist are called to help make a dairy farm more efficient. The biologist tells the farmer that he should feed the cows certain hormones to make them lactate more. The farmer asks how much it'll cost and the biologist says it'll cost many thousands of dollars. The engineer proposes to make a better milking machine to get more milk per cow. The farmer asks how much it'll cost and the engineer says it'll cost many thousands of dollars. The farmer then asks the physicist how much his idea will cost. The physicist says "It'll cost nothing and can be implemented immediately!" The farmer was astonished and ask how this is possible. The physicist responds, "Now assume a spherical cow....."

The spherical cow idealization is different from the spherical Earth idealization: it is just fundamentally wrongheaded and does not yield anything even vaguely close to the truth. Logical omniscience is more like a spherical cow than a spherical Earth. And in many respects, using a probability function to represent credences is, at a fundamental level, more like a spherical cow than a spherical Earth.

Fortunately, Bayesianism can be formulated as part of a theory of rationality without either bad idealization, so that no "problem of old evidence" or "problem of new theories" ever arises.

Objective Chance and the Principal Principle

In the foregoing section, we slipped surreptitiously between claims about credence and claims about chance. For example, we accurately calculated that on the hypothesis that a die is fair—that the chance of a six coming on any throw is $1/6$ and the chances on successive throws are independent of each other (statistically independent, uncorrelated)—the chance of getting 20 sixes in a row is $1/3,656,158,440,062,976$. And we concluded that the *credence* one ought to have in the conditional (The die comes six 20 times in a row | The die is fair) ought to stand in a ratio to Top like this:

(The die comes six 20 times in a row | The die is fair):Top:: $1/3,656,158,440,062,976$:1

That transition of thought is so natural and unobjectionable as to be almost cognitively invisible. It is an instance of what David Lewis called *The Principal Principle*.

There has been a tremendous literature about possible shortcomings and revisions of Lewis's principle in light of reliable clairvoyants, knowledge of past events, etc. This essay is long enough already without taking all of those issues on, so I am going to just bracket them. *At least in a wide variety of circumstances* the Principal Principle is unobjectionable: the degree of credence one ought to have in an event conditional on a claim which entails its objective chance is—or better is proportional to—that objective chance. The more objectively likely a proposition is according to a chance hypothesis, the higher its credibility conditional on that hypothesis. That principle provides a very useful conceptual link between credence and chance.

What I want to investigate instead is the question of the structure of objective chance itself. I have already argued for many theses about credence: 1) Rational credence obeys the Euclidean Principle and (therefore) 2) The comparative and ratio structure among credences is not always Archimedean and (therefore) 3) Credences cannot be accurately represented by the real numbers. I have also contended that 4) Sometimes rational credences are incomparable, i.e. neither $Cr(P) > Cr(Q)$ nor $Cr(Q) > Cr(P)$ nor $Cr(Q) = Cr(P)$, and (therefore) 5) The comparative structure of credences cannot be represented by assigning numbers (of any sort!) to propositions and reading off the relative credence of the propositions from the relative magnitude of the numbers. Credences and numbers are not a good fit.

But according to the Principal Principle, credences and objective chances are—at least in many circumstances—a good fit. The relations among one’s rational credences ought to mirror the relations among the objective chances. And putting these two theses together we get the somewhat surprising conclusion that numbers and objective chances are not a good fit.

That is precisely what I want to argue now.

Much of the argument has already been made. I said at the outset that for anyone who understands the concept, objective chances must be regarded as obeying the Euclidean principle. In particular, if there is some chance of P happening and some chance of Q happening and some chance of Q happening without P (i.e. some chance of $(Q \ \& \ \sim P)$ happening), then the chance of $(P \vee Q)$ happening is strictly greater than the chance of P happening: $\text{Ch}(P \vee Q) > \text{Ch}(P)$. That seems analytic: if you deny that, then I really have no idea what you are taking about. And in this respect talk of rational credences about where the dart will land on the dartboard and talk of the objective chance of the dart landing somewhere on the dartboard are effectively isomorphic. The relations of $>$, $<$, and \approx among the rational credences *just are* the relations of $>$, $<$ and $=$ among the corresponding chances. Indeed, we *read off* the relations of the rational credences from the relations of the objective chances via the Principal Principle. So if the structure of the credences cannot be captured by any function from propositions into the reals, neither can the structure of objective chances. QED.

I imagine that the rejection of numbers—any numbers, including hyperreals—as adequate representations of objective chances will be regarded as a much more controversial thesis than the rejection of numbers as adequate representations of credences. For example, it is quite intuitive that some credences are incomparable in strength, such as our old friends (for me in 2020) $\text{Cr}(\text{The Patriots win the Super Bowl in 2030})$ and $\text{Cr}(\text{Argentina wins the World Cup in 2030})$. Relative credences, it seems, are the sorts of things apt to be “fuzzy” in this way, and presuming they are perfectly sharp immediately strikes one as an unrealistic idealization, more cow than Earth. The human mind and its cognitive states are just mushy.

Objective chances, in contrast, intuitively strike one as sharp sorts of things. If there are non-trivial objective chances at all (and one might dispute that there are), every objective chance ought to be definitely comparable to every other one. With regard to objective chance, at least, we ought to have exactly one of $\text{Ch}(P) > \text{Ch}(Q)$, $\text{Ch}(P) < \text{Ch}(Q)$ and $\text{Ch}(P) = \text{Ch}(Q)$ for every pair of propositions P and Q ascribed an objective chance.

I am going to reject this claim too, and give examples to refute it.

Now at this point the reader, if the reader has been paying attention and following, ought to be getting pretty upset. For we already have a universally accepted axiomatization of probability theory provided by Kolmogorov in 1933, and what I just asserted violates his axioms. Yes, it does. I will be arguing against Kolmogorov as well. If that means that the reader's subjective credence that my position is correct is essentially infinitesimal, so be it. As long as it is not Bottom. But with the stakes plainly on the table, let's proceed.

First, let's recall Kolmogorov's axioms. This presentation is drawn from Alan Hájek's article "Interpretations of Probability" in the *Stanford Encyclopedia of Philosophy*:

More formally, let Ω be a non-empty set ('the universal set').

A *field* (or *algebra*) on Ω is a set \mathbf{F} of subsets of Ω that has Ω as a member, and that is closed under complementation (with respect to Ω) and union. Let P be a function from \mathbf{F} to the real numbers obeying:

1. (Non-negativity) $P(A) \geq 0$, for all $A \in \mathbf{F}$.
2. (Normalization) $P(\Omega) = 1$.
3. (Finite additivity) $P(A \cup B) = P(A) + P(B)$ for all $A, B \in \mathbf{F}$ such that $A \cap B = \emptyset$.

Call P a *probability function*, and (Ω, \mathbf{F}, P) a *probability space*. This is Kolmogorov's "elementary theory of probability".

Right off the bat we can see a tension: Kolmogorov requires that probabilities be represented by real numbers. We have already bought into an Archimedean structure, before we even get to Axiom 1! But there are more surprises.

Kolmogorov has not merely a domain of objects Ω , he has a (possibly special) set of subsets of Ω , given by \mathbf{F} . \mathbf{F} is required to be a σ -algebra, so it will have some convenient properties: it is closed under union, intersection and complementation, and it must include Ω itself. But \mathbf{F} is *not* required to be the power set of Ω , i.e. the set of all subsets of Ω . This is rather a key aspect of Kolmogorov's approach. Any subset of Ω not in \mathbf{F} is denominated

an “unmeasurable set”, and the function P *does not* assign any real number—any “chance” or “probability”—to it. Many paradoxes (such as the Banach-Tarski paradox which we will discuss anon) are supposed to be “resolved” by the declaration that they make use of unmeasurable sets.

What one would have naturally expected is that probability will be defined for *all* the subsets of Ω , not just some privileged few. And indeed, the Euclidean principle demands as much. Let Ω be the set of all propositions that assert that the dart will land on a particular point on the dartboard. For every subset of Ω there is the corresponding proposition that the dart lands in that subset. Then how can the notion of probability or relative likelihood *not* be defined over the power set? For every subset there is the corresponding proposition, and every proposition is strictly more likely or probable than any proposition corresponding to a proper subset, and strictly less likely or probable than any proposition corresponding to a proper superset. That follows from the postulate that every point on the board might be hit. Restricting \mathbf{F} to less than the power set makes it unsuitable for the field over which chances are defined. But if there is no restriction, then there is no need to even mention \mathbf{F} in the axioms. Wherever \mathbf{F} occurs, just replace it with the power set of Ω . Done.

What of the Axioms? Well, if you are going to use numbers to represent chances, then it seems sensible to restrict yourself to non-negative negative numbers. What would a negative number even purport to represent? Axiom 2 requires $P(\Omega) = 1$. If we postulate that the dart must hit somewhere, then the likelihood or chance of Ω is as high as chances or likelihoods can get. In our terms, that means that $\text{Ch}(\Omega) = \text{Top}$. Axiom 2 requires associating the Top chance with the real number 1. Fair enough. Top is the “multiplicative” identity in our system, so the natural way to represent it is by assigning it the multiplicative identity in the numerical representation. Finally, there is Axiom 3: Finite additivity. Well, that’s fine as far as it goes, but it doesn’t go far enough. Chances can be ascribed to the disjunction of infinite sets of propositions, not just finite ones. That should be reflected in some analogous operation on the representations, if you insist on a representation using numbers.

But now, if you are Kolmogorov, it’s all over. If you want to use real numbers your representational medium then you are out of luck. They are just not up to the task, as the

ratios between the integers are not up to the task or representing the ratios of lengths of line segments on the Euclidean plane. If one insists on using the power set as \mathbf{F} , and on the Euclidean Principle, and on using the real numbers as the range of the function P , then you just can't succeed. No "probability function" will have all the properties it should have to accurately represent probability or chance.

But once again: so what? If you can't employ real numbers—or any numbers—to represent objective chances or likelihoods, then you can't. Find some other way to do it.

And the right way it already to hand: use exactly the same structure we have already introduced for credences. In the case of the dartboard, for example, to every subset of points on the board there corresponds the proposition that the dart lands in that set. Now we want to assign *probabilities* or *chances* to that class of propositions, chances that obey the Euclidean principle. Well, they are going to look *exactly like* the way we assigned rational credences in these cases! There is a relation of comparative credence that obeys the Euclidean Principle. Some propositions are exactly as credible as others, some are strictly more credible than others, some incomparable in credence to others. In parallel fashion, some propositions are exactly as probable as others, some strictly more probable than others, some incomparable in probability to others. By building the structure of chance and of credence out of the same materials, the Principal Principle becomes trivial: match your relative credence exactly to the relative chances.¹² If you feel certain that a situation is governed by an objective chance, and the relative chances have a specific structure, then let your relative credences have the same structure.

The idea that two propositions may each have some objective chance of being true and yet those chances not be comparable to one another—neither is one greater than the other nor do they have the same chance—is certainly disconcerting. If you have become too inured to thinking of chances in terms of real numbers, then it may even sound impossible. But of course it is not "impossible" in the sense that one cannot specify a theory

¹² Again: there are issues surrounding the problem of "inadmissible information" with regard to the Principal Principle. Here, I wish to just bracket that issue: as far as I can tell, it has nothing at all to do with the issue of realizing the Euclidean Principle or the other issues I am focused on.

of chance with this feature: I just did. Furthermore, I do not think that the existence of incomparable chances is a *bug* in the program that needs to be fixed, or discounted, or minimized or hidden. It is instead a *feature* of the theory: we *want* to have incomparable chances. We *need* to have incomparable chances. Incomparable chances are exactly what *resolve* paradoxes rather than *create* them. This topic is important enough to deserve its own section.

Objections to the Euclidean Principle and the Resolution of Paradox

Many philosophers have been attracted to the Euclidean Principle and its close cousin Regularity. But it has been recognized that there is a price to pay if one wants to maintain the Euclidean principle, with disagreement about whether, and how, to pay that price. A few examples can serve to illustrate the situation.

Vieri Benci, Leon Horsten, and Sylvia Wenmackers (2018) point out that in some circumstances it is impossible to maintain these two principles:

Euclidean Principle (EP): If $A \subset B$, then $s(A) < s(B)$.

Humean Principle (HP): If the elements of A can be put in a one-to-one correspondence with the elements of B , then $s(A) = s(B)$,¹³

where “ s ” represents a rule that assigns numbers from some field to subsets of a set. I’m not sure where Hume ever announced such a principle, but the existence of a one-to-one mapping (and indeed of a much more extensive isomorphism) between objects will play a central role in the arguments we are about to consider. The general form of the Humean principle is that any two appropriately isomorphic objects must be regarded as “the same” or “equal” in some respect. Benci et. al. see that the Euclidean Principle cannot be upheld in some circumstances if s assigns members of any Archimedean field to certain propositions, and so advocate for a non-

¹³ Benci et. al. p. 522.

Archimedean field. But that alone is not always enough to salvage the Humean Principle. At the end, one must choose between the two (or else embrace strict finitism, as we will see).

Alan Hájek's paper "Staying Regular?" is (unsurprisingly) an investigation of the cost of defending the regularity principle—that anything possible should have a chance greater than zero—which is one instance of the Euclidean Principle. Ultimately Hájek argues that the price to save regularity is too high, although he does not exactly put it in that way. He rather suggests that regularity is just *ruled out* by various considerations: as he says that the end of the paper: "The trouble is that regularity appears to be untenable". Some of the untenability is supposed to arise because if one demands that *all* probabilities be represented by elements of *the same* field, pumping up the cardinality of the propositions to be assigned probabilities will eventually overwhelm the possibility of regularity. But that is not the argument I want to focus on here. The argument of interest is one that appeals to a version of the Humean Principle. Hájek describes the situation this way: "The most difficult, but also the most technically rigorous example, is one in which X is a non-measurable set – a set that simply cannot be assigned a probability, consistent with certain symmetry constraints that are forced upon the agent". What exactly are these "symmetry constraints" that are "forced on the agent", and why do they cause trouble?

Functions of any sort can display symmetries. Suppose the function f is defined over the elements of some set Ω . And suppose there is a mapping G of elements of Ω to elements of Ω . Then we can say that f is symmetric with respect to G , or f respects G as a symmetry, if for all $E \in \Omega$, $f(E) = f(G(E))$. Since we are interested in assigning chances (or credences) to elements of a set of propositions, let Ω be that set and f be either the chance function Ch or the credence function Cr . When Hájek says that "certain symmetry constraints are forced on the agent" he means that for some such function on propositions G , any chance function must have, for any proposition P , that $Ch(P) = Ch(G(P))$ and for any rational agent $Cr(P) = Cr(G(P))$.

The particular symmetry Hájek has in mind is *translation symmetry* or the more general symmetry under *rigid motions*. And indeed, if one demands such a symmetry then regularity and the Euclidean Principle are in trouble. The examples of this are familiar, but worthy of close consideration.

Hájek avails himself of our old standby example: throwing darts at a dartboard with equal chance of hitting any point (and therefore equal rational credence that any point will be hit). There also will be some chance of hitting any collection of points. Now the symmetry constraint under rigid motions (including translations) is this: the chance of hitting any set of points must be equal to the chance of hitting the result of a rigid translation of that set of points. The rigid motion must keep all of the relative distances the same, and hence is an isomorphism of the geometrical structure of the set. As a slogan: you can't change the chance of hitting a set of points on the dartboard just by moving it around, without changing its shape.

That principle is *prima facie* intuitively compelling. Certainly, nothing in everyday life would contradict it. But it is incompatible with the Euclidean Principle (and with Regularity), so if you insist on respecting the symmetry the Euclidean Principle must go. Let's see why.

Consider once again our dart board. Draw the vertical radius from the center to 12:00, omitting the center point. Even though it is infinitely thin and has no area, there is some chance that the dart will hit that radius. Now move one radius distance along the circumference, locate that point, and draw the radius from there. There should be, intuitively, twice the chance that the dart hits one or the other of these than that it hits either particular one. Rinse and repeat: keep moving one radius along the circumference and drawing in radii. Since π is irrational, this procedure will never arrive at a point that has already been chosen. So there is a denumerably infinite set of such radii. Union them all together. Call that set of points Θ . Θ is a specific region on the dartboard, and the thrown dart has some chance of hitting Θ . Hitting it is not impossible. Whatever that chance is—however one represents it, using real numbers, hyperreals, or whatever—call that chance $\text{Ch}(\Theta)$.

Now, take the region Θ and subject it to a rigid motion. In particular, rotate it one unit along the circumference, so the vertical radius at 12:00 rotates into the radius one unit along, the radius one unit along rotates into the one two units along, etc. Call the rotated image of Θ Θ' . The problem, of course, is that $\Theta' \subset \Theta$, but at the same time Θ' is a rigid rotation of Θ . And—to bring in the Humean Principle—there is an obvious one-to-one structure-preserving map from Θ' to Θ . So by the Humean Principle or by Hájek's symmetry principle we must have $\text{Ch}(\Theta') = \text{Ch}(\Theta)$, but by the Euclidean Principle we must have $\text{Ch}(\Theta') < \text{Ch}(\Theta)$. Contradiction. They can't both be true. Appeals

to non-Archimedean fields or sets of functions or whatnot are of no avail. The Euclidean Principle flatly contradicts the Humean Principle *qua* symmetry principle in this case.¹⁴ Something's gotta give.

Indeed. But the obvious thing to give is the Humean Principle in the guise of the Symmetry Principle! This is a case—which can only arise when some sort of infinity is afoot—where an object can be isomorphic to its own proper part. Such cases are already extremely counterintuitive and surprising because no familiar object of experience behaves that way. Just in the same way—because it is at root the same phenomenon—Hilbert's Hotel is paradoxical and surprising. In any finite hotel, if every room is occupied you have no choice but to turn away a new customer: there is no room at the inn. But in Hilbert's infinite hotel, the new customer can be accommodated by having everyone move down a room (a "rigid motion"). Ok: that is surprising and counterintuitive. It is not the sort of property any actual hotelier is familiar with, even if she runs a *really big* hotel. Anyone, on first hearing the example, has every right to be surprised and amazed. But it is what it is, and has nothing in particular to do with either chance or rational credence. If an object can be isomorphic to a proper part of itself—and can (therefore) be rigidly moved into a proper part of itself—it can. Accept it. Live with it. And if the perfectly innocent Euclidean principle then implies that the chance of a dart hitting a region can be strictly smaller than the dart hitting the image of the region under a rigid motion, well that's just what's to be expected. The paradox was already there before any issues of chance or credence arose.

Rotating Θ into Θ' is just an intuitively surprising sort of thing. Trying to leverage that surprise into an argument against the Euclidean Principle is not fair play.

And once we see that we see that violations of the Humean Principle and of the symmetry principle are a dime a dozen, so long as we have something infinite to play with. Take an infinite Euclidean plane with a single (complete, inextendable) line in it. Use the line to divide the plane into two regions, called "Right" and "Left". If it makes sense to throw a dart at the plane so that every point has an equal chance of being hit—as we have been wont to suppose—then there is some chance that it hits in Right. Now move the line parallel to itself to the right, so it the part to

¹⁴ One can get out of the problem by denying that the situation is possible, if one is a strict finitist. We will discuss this anon.

the right of the new line (call it Right') is a proper subset of Right. Obviously, the chance of picking a point in Right' is less than in Right, since there is a whole infinite strip that the dart can land in which is in Right but not in Right', but every point in Right' is in Right. So that simple case already violates the Humean Principle and the symmetry principle. So much the worse for them.

Can we quantify the chance of the dart landing in Right and the chance of it landing in Right' (and in all the regions of the plane) by using real numbers—or any numbers? Nope. So much the worse for that.

The case of Θ and its image under a rigid motion, though, illustrates a more important point. The Euclidean Principle renders a strict verdict about Θ and Θ' : since $\Theta' \subset \Theta$ and the dart could land on a point in Θ but not in Θ' , $\text{Ch}(\Theta') < \text{Ch}(\Theta)$. But what if we only rotate Θ by half a radius, or if we just translate it to the right, or do any other of the infinitely possible rigid motions. What then?

Move Θ under any rigid motion so its image is Θ'' . According to the Humean principle and to the symmetry principle, the chance of the dart hitting Θ must equal the chance of it hitting Θ'' . According to the Euclidean Principle the chance of hitting Θ'' is less than the chance of hitting Θ if $\Theta'' \subset \Theta$ and greater if $\Theta \subset \Theta''$. But what if neither is the case? Then as far as the Euclidean Principle is concerned, the chances are incomparable: $\text{Ch}(\Theta) >< \text{Ch}(\Theta'')$. We have already met incomparabilities in the theory of rational credence, of course, but one might have suspected that they only arise via the fuzziness of human thought. But that is not at all the case. There are also incomparable objective chances, at least if it makes sense to say that any of an infinite set can be chosen and has an equal chance of being chosen.

At this point, the reader should be getting uncomfortable. Accepting incomparable objective chances may already be a bit of a shock. But it is not just that: one might suspect that these incomparable chances will completely overwhelm the whole system. After all, if one specifies two subsets of the points on the dartboard, typically neither will be a subset of the other: there will be some points in one but not the other and vice versa. But if the Euclidean Principle only renders verdicts of comparable chance when the subset relation obtains, then it will be almost completely useless. This—unlike the conflict with the Humean Principle or with symmetry—is indeed a serious objection.

Fortunately, it can be easily overcome.

The Return of Kolmogorov

Above I launched an argumentative campaign against taking Kolmogorov's axioms for probability functions as a *foundational ground* for any account of objective chance or of rational credence. I stand by those arguments. But there is nothing at all preventing the use of Kolmogorov functions (as we may call them) as an essential *component* of a theory of objective chance. This also allows us to accept all of the work that has been done in the Kolmogorovian tradition with only a very, very mild adjustment. It is not that one can have one's cake and eat it too—that is obviously impossible—but one can have a Euclidean cake slathered with as much Kolmogorovian icing as one likes. Let me explain.

The problem with Kolmogorovian accounts of chance or probability—presented as foundational—is essentially the problem of zero probabilities, and therefore also the problem of equal probabilities (i.e. probabilities whose difference is zero). If the Kolmogorov function were all there were to say about the matter, then all “zero probability” events would have to be regarded as exactly the same with respect to chance, and all “equal probability” events also regarded as the same. But they aren't. The dart landing at p and the dart landing at either p or q are both zero probability events in the usual telling, but the latter is more likely than the former. The obvious solution is to *ignore the Kolmogorov function when it returns a value of zero*. Two events having the same Kolmogorov “probability” means nothing: they could, consistent with that, have the same chance, or one could have a higher chance than the other, or their chances could be incomparable. However, when the Kolmogorov “probabilities” are *different*, when one is greater than the other, then take that seriously. Call a Kolmogorov “probability” function K . If we take K seriously, then if $K(P) > K(Q)$ we set $\text{Ch}(P) > \text{Ch}(Q)$ and $\text{Cr}(P) > \text{Cr}(Q)$. But if $K(P) = K(Q)$, we just ignore it and look around for some other principle that would settle the matter. Sometimes the Euclidean Principle will.

Let's work through an example. The real numbers provide an infinite universe, with all of its subsets, and the elements ordered by the relation \geq . There was long a question of how one might

measure sets of real numbers. Of course, each set has a cardinality, but that is not discriminating enough for the purposes desired. So Lebesgue came up with an obvious solution.

First, for any closed interval of the reals—any set consisting of all the reals between p and q inclusive—the obvious measure of that set would be $|p - q|$. Thus there are “twice as many” real numbers (by this measure) between 14 and 16 than there are between 9.6 and 10.6, even though the sets have the same cardinality. Any intuitively acceptable “measure” of the “real number line” should have this property. It follows that a single point p , considered as an interval both of whose end points is p , must have zero Lebesgue measure.

Well, what else? Another natural thought about a measure is that it should be additive for two disjoint sets: if $\Sigma \cup \Sigma' = \emptyset$, then $\mu(\Sigma \cup \Sigma') = \mu(\Sigma) + \mu(\Sigma')$. So Lebesgue demanded this as well. Not only does μ now provide a measure of all intervals, it provides a measure of all finite sets of intervals (again, including points as degenerate intervals). Now one might like to make the additive property more extensive: not only is the measure additive for pairs of disjoint sets, it is additive for all sets of disjoint sets, even infinite sets. But that is a bridge too far: since every subset of the reals is the union of a set of disjoint sets of measure zero—namely the singleton sets of each point in the subset—the whole thing threatens to collapse into triviality or contradiction. It turns out that one can do a bit better than finite additivity: one can allow countable additivity without problems, because no non-trivial interval can be composed of countably many points. (If it could that would mean trouble, because the interval would have to have positive measure by the fundamental postulate but also zero measure if the “sum” of any collections of zeros must be zero.)

There is one more important feature of Lebesgue’s measure over the reals: the field over which it is defined is a sigma algebra, which means that not merely (in this case countable) unions and intersections of Lebesgue-measurable sets are Lebesgue measurable, but complements are as well. If the measure of a set is μ , then the measure of the complement is $1 - \mu$. This extends the reach of the measure beyond countable sets of intervals. For example, the rational numbers are a set of Lebesgue measure 0 in the unit interval, since it is a countable set of points, each of which has measure 0. So the non-rational reals between 0 and 1 form a set of Lebesgue measure 1, even though it contains no finite intervals at all. In that sense, “almost all” real numbers between 0 and 1 are—according to Lebesgue measure—irrational.

Lebesgue measure obviously does not satisfy the Euclidean Principle because the measure of the set $\{p\}$ is identical to the measure of $\{p, q\}$ (viz zero) even though $\{p\} \subset \{p, q\}$. The

Lebesgue measure of a closed interval is the same as the corresponding open and half-open intervals, because they only differ by a set of measure zero. So the Lebesgue measure is not the right beast to use for specifying the chance of a randomly picked real falling in a certain set of reals, or of a randomly picked real between zero and one falling in a set. And that's because it is not the right foundational measure for judging exact equality of areas or regions.

However, there is nothing at all wrong—and everything right!—in insisting that the chance of the number falling in Σ is greater than it falling in Σ' whenever $\mu(\Sigma) > \mu(\Sigma')$. That extends the comparability of objective chances immensely, and in exactly the way we want. We can now say that the chance of the randomly chosen number falling between 0 and 1/4 is less than it falling between 1/2 and 1, even though neither interval is a proper part of the other.

Appending the Lebesgue measure in this way to the Euclidean Principle yields—I claim—a much improved theory of objective chance for the “randomly thrown dart” or the “randomly chosen number”. If we add that each number is equally likely to be chosen then we can calculate ratios of chances for all sorts of events. The chance of the randomly chosen number being either 1/4 or 1/2 is twice as big as it being 1/2, or indeed of it being any other particular number. Of course wedding Lebesgue and Euclid requires slightly neutering Lebesgue—we no longer pay attention to the Lebesgue “measure zero” judgements—but the union returns every intuitive judgment of relative chance and ratios of chance we would want.

There are, however, limitations to Lebesgue just as there are to Euclid. Not every pair of sets stands in the subset relation, which puts a severe hamper on the scope of the Euclidean Principle. And we have already limited Lebesgue by ignoring judgments of zero or equal probability. But there is another limitation to Lebesgue: some sets of points fall outside the scope of his measure entirely. These sets are just not “Lebesgue measurable”. We don't run across such sets in everyday life, or any applications of mathematics to real physics, so they shouldn't bother us too much. But clearly, if a set is not Lebesgue measurable then *Lebesgue* will be of no direct help in assigning a chance that a random real be a member of it.

It is possible that Euclid can be of some assistance here. A Lebesgue unmeasurable set can have a Lebesgue measurable subset. Indeed, the union of any measurable set with any disjoint unmeasurable set is unmeasurable. But Euclid assures us that the chance of the chosen number being in the union is greater than it being in the measurable subset, so that helps. And of course Euclid tells us that the chance of the randomly chosen number being in the unmeasurable set is

greater than it being in any proper subset. So the team of Euclid + Lebesgue is much more powerful in these circumstances than Lebesgue alone or Euclid alone. Nonetheless, there are still cases where no judgement of $>$, $<$, or $=$ is rendered about the relative chance of the real number being in one of two sets. In that case, we are left with $\text{Ch}(\Sigma) \succ \text{Ch}(\Sigma')$ and $\text{Cr}(\Sigma) \succ \text{Cr}(\Sigma')$.

Similarly, any Kolmogorov probability function that has been found useful in physics can be taken on board—properly neutered—and used in conjunction with Euclid. The only possible applications that could be affected would be ones that made a particular sort of use of the equality of the Kolmogorov function, rather than the inequality of the function. Any such case should be carefully examined because the reasoning could be faulty. I personally doubt that any of the actual uses in physics are like that.

Finally, we can say a word about the Banach-Tarski Paradox. That paradox has gotten quite a lot of attention from philosophers because the result is both extremely counterintuitive and is proved using the Axiom of Choice. Many have taken this as grounds to question that axiom.

The proof requires three dimensions, so we need a measure over a three-dimensional space—in particular the flat Euclidean three-dimensional space E^3 . There is a natural extension of Lebesgue measure from R^1 to R^3 , and then one can use elements of R^3 as coordinates on E^3 via the Cartesian method in the usual way. So we can speak loosely of “Lebesgue measure” on E^3 , although properly that is a misnomer. In any case, what Banach and Tarski prove (and I say “prove” because the Axiom of Choice is true) is that one can take a solid unit sphere, partition it into five pieces, move the pieces around rigidly in the space, and then “reassemble” them into 2 solid unit spheres. That is certainly surprising behavior. It will come as no shock that the pieces are not Lebesgue measurable. If they were, the volume of the union would be fixed irrespective of their relative locations (so long as they don’t overlap), because the Lebesgue measure *is* invariant under rigid motions. (This is obvious because an interval always rigidly moves into another interval of the same length.) The fact that the Lebesgue measure is invariant under translations, of course, has no bearing on rotating Θ since the difference in Lebesgue measure between Θ and Θ' is zero, which has no significance.

So the Banach-Tarski proof involves several essential components, some of which are somewhat unfamiliar. It involves Lebesgue-unmeasurable sets, which we never run into in everyday life. It involves appeal to the Axiom of Choice, which we also never do in everyday life. It involves the tacit assumption that “things don’t change size under rigid motions” which we *do*

employ *all the time* in everyday life. And it essentially involves infinite sets of points, which we don't deal with—at least not as such—much in everyday life. So when such a surprising result falls out, we naturally look for a perpetrator to blame. Both the use of unmeasurable sets and the appeal to the Axiom of Choice have been fingered as prime suspects.

The main suspects, of course, naturally try to cast blame on one another. Defenders of the Axiom of Choice can say that there is nothing wrong with the Axiom—or the proof—and the counter-intuitiveness arises from the use of unmeasurable sets. Critics of the Axiom want to deny that response, reject the Axiom and maybe unmeasurable sets as well. But I think what we can see is that all of this is just a distraction. The real problem lies neither with the Axiom of Choice nor with unmeasurable sets. The real problem was before our eyes long before any of those issues came up.

Go back to our fan of radii Θ , with the radii shifted by one radius length along the circumference started from vertical. Θ can be rigidly rotated into a proper subset of itself—it can be made to “shrink” by a rigid rotation. It can also be rigidly rotated into a proper superset of itself: just twist the other direction. All of these sets have Lebesgue measure zero—which does not change under the motion—but so what? Lebesgue measure zero means nothing, and the Euclidean Principle proves that indeed Θ can become larger or smaller just by a rigid rotation. *And this result makes no mention of the Axiom of Choice or of any unmeasurable set.* It is just a property of certain rigid motions in a continuum. Like or lump it, that's what it is, and the Banach-Tarski situation is in principle no better and no worse. Once you have accepted the something can change size by a rigid motion—which common sense rejects on account of everyday experience—you have already made the leap. Banach-Tarski just shows how big the leap is once you get to three dimensions.

So if one is unwilling to accept Banach-Tarski, then rejecting the Axiom of Choice or rejecting Lebesgue-unmeasurable sets really is not enough. You also ought to reject the possibility of the radii fan Θ rotating into a subset of itself. And the only principled way to do that, I think, is just to reject real infinity altogether. One could hold that no really existing thing—or even abstract thing perhaps—can have infinitely many parts. This rejects, at one fell swoop, both infinite divisibility (as in a continuum) and infinite extension. If you do that, then all the problems with Banach-Tarski and with the radii fan and many others as well simply evaporate. These sorts of problems do not arise in strictly finite settings.

The paradoxes arise from the circumstance that there are things that are completely isomorphic to their own proper parts. And that feature is always an indication of some sort of infinity. Galileo, in *The Two New Sciences*, already argued for the inherently paradoxical—or perhaps incoherent—nature of infinity by pointing out the obvious one-to-one mapping from the integers to the squares. In one sense (Euclid’s) there are obviously fewer squares than integers, and indeed the squares become unboundedly sparser and sparser as you go to higher and higher integers. But on the other hand, the relation shows that the integers and the squares can pair up with nothing left unpaired. That—in a finite setting—means there are the same number of items. Cantor, of course, went with the one-to-one criterion and fashioned a perfectly acceptable theory of cardinality out of it. But we should never forget that infinity is, as it were, shot through with paradox from beginning to end. When we get unsettling results in infinite domains the source of paradox can just be the infinitude itself, not the means of deriving the result. And one might therefor reject infinity root and branch.

I think this is a reasonable position to advocate. If it is correct, then physical space or space-time cannot be infinitely divisible: it must be ultimately discrete. I have myself developed a geometry for fundamentally discrete spaces and space-times in *New Foundations for Physical Geometry: Full Discrete Geometry*. And in a completely discrete setting, most of the paradoxes concerning objective chance also go away. But the proper thing to do is to address the fundamental problem of the reality of infinity all on its own, not in a setting like this.

Probability Densities

Kolmogorov was concerned with probabilities. Measure theory is concerned with measures more generally, including measures that are not finite and hence not normalized or normalizable. For example, Lebesgue measure over the entire real line is unbounded in extent since there are intervals of greater length than any specified length. So although Lebesgue measure over the unit interval is formally a Kolmogorov probability measure, Lebesgue measure over the entirety of the reals is not. Nor is the standard measure of any dimensional Euclidean space.

Physical probabilities, though, are often inextricably related to geometrical measure. In our standard sort of dart-throwing example, where the dart is specified as “equally likely to land on

any particular point”, our judgments about chance are parasitic on our judgments about geometrical size. Indeed, the way the physical chances are *specified* is mathematically parasitic on measures of geometrical magnitude that have nothing *per se* to do with chance at all.

This is manifested in the fact that what is commonly used in these cases is not a probability measure at all: it is a probability *density*. That is, one is given a function f that has to be integrated over some geometrical domain in order to yield a physical chance. What is commonly called a “flat” or “uniform probability measure” is really a constant function that is supposed to represent a probability *density*. In order to get a probability, one must integrate that function over a measurable geometrical domain. Geometrically unmeasurable domains, of course, are assigned no probability at all. And domains of geometrical measure zero—such as individual points or denumerable collections of points—all automatically get assigned probability zero because the domain is assigned geometrical size zero.¹⁵

There is no reason that a Kolmogorov probability function must be specified this way, i.e. using a probability density meant to be integrated over a domain with some other metric already defined on it. But where things are done that way—as is common in physics—the density provides us more resources to specify physical chances.

Let’s take the example of our dartboard yet again. We can posit a process that picks a point in the dartboard in such a way that “every point has the same chance of being picked”. In such a case, we obviously use a constant probability density function. Or, we could imagine a process that could pick any point, but favors the points closer to the bull’s-eye, as a skilled dart-thrower would. For example, if the dartboard is covered by polar coordinates and has unit radius in those coordinates, we might specify the probability density $f(r,\theta) = \frac{3}{\pi}(1 - r)$. (The $\frac{3}{\pi}$ is a normalizing factor, so the density integrates to 1 over the entire board.) Note that the probability density is a function defined *at every individual point*, and which is different at different points, even though the *probability* it integrates to at every point is the same—viz. 0—because the geometrical measure of every point is 0.

Given this probability density, we would not be inclined to say that every point has the same chance of being hit, even though the Kolmogorov probability function for each point is the

¹⁵ If one wants to avoid this result then one can employ not a probability density *function* but rather a *distribution*, such as the poorly-named “Dirac Delta Function” (which is not a function at all). This section is about the use of proper probability density functions.

same. We would rightly say that the closer a point is to the center, the more likely it is to be hit. Indeed, we would say that for individual points *the chance of it being hit stands in the same ratio to the chance of another point being hit as the probability densities at the point stand to each other*. For example, any point half-way from the center to the edge has half the chance of being hit as the center does.

As with the Kolmogorov probabilities, positive differences between the probability densities at points must be taken seriously as reflecting the relative chances of those points being hit. Whether we must regard points with the same density as being exactly as likely to be hit is a different matter. Off-hand, I can see no objection to doing so, and that certainly seems to be the principle that underlies our intuitive judgments in these cases. Further, I can see no objection to taking the ratios of magnitudes of the probability densities at different points (if they exist) as exactly the ratio of the likelihoods that the points be hit. In this way, the implications of the probability density for the objective chances is extended beyond the implications that derive from the probabilities (which, again, are all 0 for individual points).

However, we must be judicious here. At this juncture, we have a set of principles for objective chances in cases where a Kolmogorov measure is specified via a probability density function that is to be integrated against a geometrical measure. The principles are:

- 1) If the Kolmogorov measure of region A is greater than that of region B, then $Ch(A) > Ch(B)$, where ‘Ch(x)’ represents the objective chance of a point in x being hit.
- 2) If the Kolmogorov measures of the regions are the same, then nothing immediately follows: the chances might be the equal, or one might be larger than the other, or they might be incomparable
- 3) If the regions are points, then the ratio of their chances ought to be proportional to the ratio of the densities, if there is one.

This allows us to make exact sense of the claim that “all the points are equally likely to be hit” or that “points nearer the center are more likely to be hit than points farther away”. And of course the Euclidean principle extends these judgments when $A \subset B$. But none of these principles settles certain cases where we have strong intuitions. For example, in the dartboard example, with either of the probability measures mentioned above, we are very tempted to say that the chance of the dart hitting any radius is the same as it hitting any other. But such a result does not follow from

our principles since the Kolmogorov probabilities are the same (indeed, both 0) and the regions are not points, so the density does not settle it.

What we want to do is to appeal to rotational symmetry. But as the example of the radius fan Θ demonstrates, unchecked symmetry arguments will run afoul of the Euclidean Principle, which must be held sacrosanct. If there is some principled way to restrict the symmetry argument so it gives the intuitive answer for the radii but not the unacceptable one for Θ , then that deserves serious consideration. But I myself have not been able to discover such a restriction.

Perhaps this is one reasonable extension that almost gets the equality of the chances to hit any particular radius right. The probability density is just a non-negative real function over the domain—in this case the dartboard. If we integrate that function against an area measure for a measurable locus of points, we can get a Kolmogorov probability, but the probability for any locus of area zero will automatically be zero. Hence the probability for any radius will be zero. But there is nothing preventing us from integrating the very same density function against a *different* spatial measure, e.g. a *length* measure of a 1-dimensional locus of points. If we do that, then we can assign numbers to all the 1-dimensional measurable sets, and then take the ratio of the chances that the dart hits in two such sets as proportional to the ratio of those numbers. This technique would almost give us the intuitively correct results: in the case of the flat probability distribution, the ratios of the chances of being hit for any two radii would be 1:1 because they have the same length, and the ratio of the chances within each radius of being within half a radius of the center to being more than half a radius would be 1:1. With the non-uniform density mentioned above the ratio chances for any pair of radii would again be 1:1, and the ratio of the chances within each radius of being within half a radius of the center to being more than half a radius would be 3:1. One should be careful not to equate a ratio of chances being 1:1 in this sense to “being equal” since the ratio of lengths of two lines may be 1:1 even though one is a proper subset of the other (by deletion of an endpoint, say). Still, a probability density has the virtue of allowing ratio comparisons of what we might call “0-dimensional chances”, “1-dimensional chances”, “2-dimensional chances”, etc. among themselves, with there being no ratios between the chances of different dimensionalities. These judgements would follow the analogous judgments we make about lengths, areas, volumes, etc.

The only caveat to bear in mind is that two chances having the ratio 1:1 is not the same conceptually as the chances being equal. As mentioned above, the ratio of the length of a closed

line segment to its open counterpart is 1:1 because the difference is of 1-dimensional measure zero. But the chance of the dart hitting the closed segment is strictly larger than hitting the open interval. Given a choice between tickets that pay the same prize if a point is hit, it is rationally mandatory to prefer the closed interval over the open. Still, there would be no amount of *money* it would be rational to pay in order to secure the closed segment over the open, since the difference is, as it were, infinitesimal.

Generic Chances and Typicality

Some chances are postulated to be dynamical chances. That is, they are written into the fundamental dynamical physical laws, which are presented as stochastic processes. In a world with fundamental dynamical chances, two isolated systems that start off in exactly the same physical state can nonetheless end up—after some period of time—in different states. The fundamental dynamical chance specifies these conditional probabilities: how likely each possible final state is given the initial one.

But even if the fundamental dynamical laws are deterministic, so that isolated systems in the same state at two times must be in the same state at all equally later times, we think that in some sense the history of the world may be understood in terms of objective physical chances. For example, in the history of the actual world, there is a precise statistic of how fair (i.e. uniformly dense and cubical within the tolerances allowed) dice thrown legally at gaming tables (so they bounce off the far wall of the table) have landed, and that statistic is that each side has come up very, very nearly one time in six. We know that because those are the odds that casinos use to predict their long-term profits and the casinos do, in fact, make the predicted long-term profits. If the statistics of the results varied appreciably from the predictions, gamblers would be able to make a profit off the casinos.

In some decent sense to be explored, a fair die (in the sense above) thrown fairly (in the sense above) by a human being has a (very nearly) $1/6$ chance of showing each of the six sides, and if such a die is thrown repeatedly the sequence of actual results given in chronological order will show about $1/6$ frequency for side and pass any test for statistical randomness in use. Or at least, if that were not true then we would both be astonished and simply have no clue what to make of it. If it were not true, that would occasion a revolutionary change in physics somehow.

All of the above claims are completely untouched if the fundamental physical dynamics is deterministic, as classical physics supposed it to be. And even if the fundamental physics is not deterministic (due, e.g., to quantum-mechanical effects), that is neither here nor there. The statistics of die throws and coin tosses are not based, in terms of the illuminating explanation, in quantum-mechanical chances. We know this for coins because a sufficiently precise coin-tossing device can be made such that a coin placed heads on the device and flipped (turning several times in the air) will always (or nearly always) land heads, and a coin placed tails will always (or nearly always) land tails. And we know this is possible because it is actual: Persi Diaconis and his collaborators have made such a device.¹⁶ In fact, as they show, the standard method for tossing a coin and catching it is not “fair”: the coin has about a .51 chance of showing the side it started out on. But we could specify a fair method of coin-flipping, such as “Have a human put the coin in a 4” cubical metal box, shake it vigorously for 30 seconds so it can be heard rattling around, then open the top”. We expect that a fair (uniformly dense flat cylindrical) coin treated in such a fashion would have a .5 chance of ending up on each side. And this would be so even if the underlying dynamics were completely deterministic.

It seems as if these sorts of chances are “objective”. Or at least we can say this: the *actual frequency of results of throws of fair dice at gaming tables* is objective—it is just a plain physical fact—and it is indeed very, very near to 1/6 for each side and randomly distributed. This is a circumstance accurately predicted by casino owners and is striking enough to demand an explanation.

It is tempting to call these sorts of chances “macrochances”, since they are specified in terms of macroscopic conditions, such as which side of the die finishes facing upwards. Or possibly “empirical chances” because they can be checked empirically through repeated experimentation. Sometimes, where the underlying fundamental dynamics is postulated to be deterministic, they have been called (slightly oxymoronically) “deterministic chances”. But the best term for them is rather “generic chances”. They are generic in the sense that the “chance set-up”, i.e. the condition upon which the conditional chances are conditioned, is only specified generically rather than in complete physical detail. For example, the description “Have a human put the coin in a 4” cubical metal box, shake it vigorously for 30 seconds so it can be heard rattling around, then open the top”

¹⁶ See P. Diaconis, S. Holmes and R. Montgomery “Dynamical Bias in the Coin Toss”, <https://statweb.stanford.edu/~susan/papers/headswithJ.pdf>.

is generic in many, many respects: the *exact* size of the box, the metal of which it is made, the particular human shaker, the exact method of shaking, the ambient gravitational field, the way the coin is initially put in the box, etc. etc. etc. are not specified. The “chance set-up” could be realized in an infinitude of physically distinct ways. And that is absolutely essential: if the chance set-up *precisely* specified the way the experiment is to be done, and if the fundamental laws were deterministic, then the conditional probabilities for the different possible (generic) outcomes would be all 1 or 0. And even if the fundamental dynamical laws were indeterministic, there is no reason to expect the coin chances to be anywhere near 50%.

Let’s take one specific and crucial use of generic probabilities in the history of physics: Maxwell’s derivation of the Maxwell/Boltzmann velocity distribution for a monatomic gas at equilibrium. The relevant probabilities are the probabilities for the atoms in the gas to have a velocity—or momentum—within a certain narrowly specified range. These are evidently not macroscopic probabilities in any sense, since the precise velocity of an atom—or the distribution of velocities in the gas—is not a macroscopic quantity. But they are generic probabilities because the state of the gas—being at macroscopic equilibrium with the macroscopic parameters of temperature and pressure macroscopically stable—is clearly a generic description that does not nail down the precise specific microstate of the gas. The fundamental dynamical laws cannot be directly applied to the generic description since they require a specific state as input, yet the generic description together with the fundamental laws somehow accounts for or explains or predicts the velocity distribution and the chances for individual atoms. Furthermore, Maxwell’s argument also yields a dynamics for the generic probability: if the velocity distribution is not of the Maxwell/Boltzmann form, it will (typically) evolve to be closer to it, at a calculable rate. The significance of the “(typically)” in the previous sentence will much concern us anon.

One must bear in mind our precise logical situation. Ultimately, the generic chances and explanations will depend on the precise underlying microdynamics. To take the obvious extreme case, if the microdynamics is static—nothing changing at all—then the generic dynamics must be as well. However, the microdynamics must be applied to a precise microstate: it simply gets no immediate mathematical purchase on a generic description. So somehow or other the generic description must be “filled in” to yield a specific one. The question is how exactly that is to be done.

Clearly, one thing that would assist in the necessary filling in is a measure (or other serviceable specification of “size”) over (at least some of) the power set of the space of possible specific states. Suppose, for example, one has an underlying deterministic fundamental dynamics. (One could equally use an underlying stochastic dynamics, *mutatis mutandis* in the obvious way. The reader can fill in the details.) That dynamics effects a one-to-one map from specific initial states to specific final states. The generic description of the chance set-up then picks out a (probably somewhat fuzzy) set of candidate precise initial states. Each of those precise initial states will be evolved by the dynamics into a precise final state, which in turn determines a generically specified outcome (such as which side of the die came up, or whether the coin landed heads or tails, or what the distribution of velocities in the gas is). The pre-image of each of the generic outcomes will be a subset of the space of possible initial conditions. And so if we had a measure of that space, we could define a measure associated with each outcome. Those measures could then be compared to each other: we might be able to say which generic outcome is more likely than which, and even define ratios between them. The question is where exactly this additional measure might *come from*.

There are three obvious suggestions, some or all of which might be appealed to in giving the answer:

- A) The measure is somehow derived from, or at least inspired by, the precise microdynamics.
- B) The measure is somehow derived from, or at least inspired by, some other measure that is already present in the fundamental physics.
- C) The measure is somehow derived from, or at least inspired by, some *subjective* or *a priori* source, such as the credences of some agent or some principle of rationality or some empirical experiences of an agent.

All of these—singly and in combination—have been put forward as the relevant additional measure needed to define generic probabilities. My own account will advert to both A) and B), but not to C). In particular, insofar as rationality principles come in at all, they come in only at the end after all the hard lifting has been done. It is *because* one can define the relevant generic chances by reference to A) and B) that they are the rational chances to use to set one’s credences via the Principal Principle.

Let's start with suggestion A. Given a space of initial conditions—that is, a space of all possible initial data sufficient for the dynamics to determine the future history of the universe¹⁷—and given the microdynamics, some measures over the initial condition space are picked out because they are *stationary*. That is, if one evolves the measure using the dynamics, letting the probability density flow along the flow lines of the dynamics in the initial condition space, the evolved measure is the same as the initial one.

Being stationary is an instance of the more general property of being *equivariant*. The property of equivariance is the property of being defined by the same rule or principle at all times. For example, in the pilot wave theory of quantum mechanics, the space of all possible initial states of the particles is configuration space: the space of all possible ways the particles can be distributed in space at a given time. Given an initial wavefunction, each configuration will evolve in a manner determined by the dynamics. Generally, no interesting probability measure over configuration space will be stationary, but the measure defined by the absolute square of the wavefunction is equivariant no matter what the initial wavefunction is: if we let the measure be carried along by the dynamics, it will always equal the absolute square of the contemporaneous wavefunction. In that sense, the absolute square of the wavefunction is a measure privileged by the dynamics itself, with no other considerations coming in.

Why should equivariance be an important feature of a measure over the initial-condition space? As we will see, we will be interested in making claims about the “vast majority” or “almost all” of the systems in a collection of systems, or of the initial states in a collection of initial states. So we will need a standard or measure by which to define “vast majority” or “almost all”. And that standard must, as a matter of conceptual coherence, be equivariant. For suppose it weren't. Then, given a deterministic fundamental dynamics, we could find ourselves obliged to say that a set of systems in a collective that constitute “overwhelmingly most” of the collective evolved deterministically by a one-to-one map into a set that is no longer “overwhelmingly most” of the collective, and that just makes no sense at all. If we want to make claims about how “most” or “nearly all” of the systems behave then we must have an equivariant standard for what counts as “most” or “nearly all”. Of course, if only finitely many initial states are possible, there might be no issue: just count the number that have the desired property and set a threshold for the proportion

¹⁷ Given a time-reversible dynamics such data will also suffice for the dynamics to determine the entire past of the universe, and so the dynamics specifies a one-to-one map from phase data at one time to phase data at any other time.

that can be considered ‘overwhelmingly most’. But if there are infinitely many possible initial conditions then we need a measure of them, and the measure must be equivariant to do the work required of it.

Condition B above adverts not to the dynamics but to other measures that are already part of the physical theory, particularly measures of space and time. In classical physics, and particularly classical mechanics, such measures are ubiquitous and often invoked without any fanfare or notice. The initial-condition space for classical mechanics is *phase space*, i.e. the space of all possible specifications of both the position and momentum of the particles in the system. Phase space comes—without adverting to any dynamics at all—with a natural measure which derives from the measures over space and time. These measures are so obvious that we use them without even noticing it.

For example, suppose that there is a single particle of mass m with kinetic energy E inside a 6” cubic box (note the generic description), and we want to calculate the chance that it collides with the walls of the box in a given second. The natural way to proceed is first to calculate its speed as $\sqrt{\frac{2E}{m}}$. If it does not collide with the walls, it will travel a distance $\sqrt{\frac{2E}{m}}$, so the question is whether a wall of the box lies within $\sqrt{\frac{2E}{m}}$ of the position of the particle in the direction of its present velocity. If the particle is not within $\sqrt{\frac{2E}{m}}$ of the walls, then it is automatically safe: it will not collide. If it is within $\sqrt{\frac{2E}{m}}$ of a wall, then the question come down to exactly in which direction it is travelling. The directions that lead to a collision will make up some solid angle that depends on the exact location of the particle. The chance of collision is then the ratio of that solid angle to the total solid angle of 4π available. By integrating over the entire initial condition space—all of the possible initial positions and velocities consistent with the generic description—one arrives at a “chance” of collision. This is a standard sort of classical-mechanical calculation.

But in order to do the relevant integration one needs a measure over the initial condition space, and one uses the obvious measure: the spatial measure for the location within the box and the measure for velocities that assigns the same chance to all directions. Formally, or course, there are an infinitude of other measures that can be defined over the initial-condition space. but it never for a moment even occurs to one to use any of those—without some reason to depart from using

the spatial measure. (What might such a reason be? Maybe a gravitational field which breaks the spatial symmetry, for example.)

The measure of space itself generates a natural measure over configuration space. Together with the measure of classical time it generates a similar natural measure over speed spaces: there are “twice as many” speeds between 1 m/s and 3 m/s than there are between 5 m/s and 6 m/s. And for the direction of a particle, we have the geometrical measures of solid angles. Since the initial condition space for classical mechanics is phase space—given by the positions and momenta of the particles—there is a natural measure defined over classical phase space by the measures of space and time themselves. And as it turns out, in classical mechanics this natural measure is equivariant: a set of systems that occupies a certain measure of phase space, by this natural measure, will evolve into a set that occupies the very same measure at all times. That is the import of Liouville’s theorem, and so this natural measure is sometime called “Liouville measure”. Liouville measure satisfies both conditions A and B above: it is easily definable from measures already in the physics and comports with the dynamics in virtue of being equivariant (indeed, static). It is really the *only* measure that anyone would ever use to characterize “the size of a region in phase space”. It is the measure that always is used, and the one any beginning physics student will use when asked a question like the one above about the chance of the particle colliding with a wall.


The natural—Liouville—measure is proportional to the geometrical measure. In some sense, that is what makes it the natural measure. But still, it is not clear exactly what proportionality really portends. As David Albert remarks, in some sense the Liouville measure seems to suggest that there should be a greater chance of finding the spatula in an apartment in the bathtub than in a kitchen drawer since the region of phase space with it in the tub is greater (by Liouville measure) than the region with it in the drawer.¹⁸ But as a matter of fact, spatulas are often in kitchen drawers and almost never in bathtubs. So that “chance” bears no relation to actually observed frequencies. And one wonders what role a chance disconnected from actual frequencies—even if it can be mathematically defined—can play in the physical explanation or even prediction of anything,

¹⁸ *Time and Chance*, p. ???

That is a slightly involved issue, which brings in more detail about dynamics than we have mentioned yet. We will approach this question via an idealized case, and then discuss the steps to de-idealize it.¹⁹

Test Case: the Infinite Quincunx

The status of generic chances can be explicated by the example of the quincunx or Galton board. That is a board with the familiar regular array of perfectly cylindrical pins in it, as depicted below:



The Quincunx

A marble is introduced through the hopper at the top, and then makes its way (guided by the exact microdynamics) down the board. For the sake of analytical simplicity (in a certain sense), we will begin by imagining the board to be infinitely long. At the first pin, the marble will pass either to the right or the left, and at each successive row the marble will pass either to the right or left of the gap it came through in the row immediately above. Therefore, each trajectory through the device

¹⁹ I learned of this approach—and the use of the quincunx as an illustration—from Detlef Dürr, Sheldon Goldstein and Nino Zanghì, but they have used it rather than explicated it for a general audience.

will generate an infinite sequence of Rs and Ls. The precise microdynamics and configuration of the board therefore maps each possible initial state at the top of the hopper into an infinite sequence of Rs and Ls.

Those sequences can then be partitioned into three classes: those in which the limiting frequency of Rs is precisely .5, those in which the limiting frequency of Rs is some real number other than .5, and those in which there is no limiting frequency at all. All of these outcomes are *possible*, in the sense that some initial conditions will yield them. Indeed, some initial condition will yield every possible limiting frequency between 0 and 1. But what we intuitively believe is that of all those possible outcomes the particular outcome of limiting frequency .5 is special, and it is the configuration of the board and the microdynamics (and not anything about anyone's credences) that make it special. We are now in a position to articulate an exact sense in which that is true.

Consider the set of all initial conditions consistent with the generic description "the ball is dropped into the hopper". The initial position of the ball must lie somewhere above the hopper, and the initial momentum be somewhere near zero (make each of these conditions precise in any reasonable way). Now considering that set of possible initial conditions, what we believe to be a mathematical fact is that a *set of measure 1* (in natural or Liouville measure) yields sequences that limit to frequency precisely .5 and sets of measure 0 yield either some other limiting frequency or no frequency at all. As Detlef Dürr has remarked, we are *infinitely far* from being able to prove that assertion. That is because the dynamics of the marble hitting and ricocheting off the pins is chaotic, so no general means of tractable exact analysis is available. Still, due to the spatial symmetry of the pins, limiting frequency of exactly .5 is the only plausible conjecture for a set with any measure other than zero. In the sequel we will simply presume this conjecture is correct. The question is how much this result can be parlayed into.

One immediate reaction is that by itself it cannot be parlayed into much. What the result implies is that in a precise mathematical sense *almost all* possible initial conditions in the set yield limiting frequency .5, where "almost all" means a set of measure one in the indicated measure. But—the objector objects—that result is only relative to the chosen measure, in this case the Liouville measure. But what, in turn, makes *that* measure special or worthy of note? The Liouville measure, for example, is proportional in its spatial aspect to the spatial measure. So if it is supposed to somehow represent the "likelihood" that the marble starts out in a particular location then it

assumes that the person releasing the marble is equally likely to release it from any particular inch-long region above the hopper. But maybe the person is right-handed, and therefore much more likely to release the marble on the right than on the left. Then the Liouville measure is *not* a good representation of where the marble is likely to begin and, as they say: Garbage In, Garbage Out. That is, conceptually the whole problem has just been pushed back to justifying the use of Liouville measure over the precise initial conditions that satisfy the generic description. Given some measure over initial conditions of course one gets a measure over the final outcomes, but all that has been done is to replace one problem with another that is equally problematic. Nothing has really been accomplished.

It is here that our idealization of the infinite board comes to the rescue. Because the board is infinite, we think that the (Liouville) measure of initial conditions that yields the precise frequency .5 is of measure 1. And if that is true, then it is also of measure 1 in *any measure that is absolutely continuous respect to with Liouville measure*, i.e. any measure that assigns zero measure to all sets whose Liouville measure is zero. So if our right-handed ball-dropper is more likely to release the ball on the right, fine! Represent the ball-dropper not by Liouville measure over the initial conditions but by some other measure that vastly favors releases on the right. Pile the measure up on the right as much as you like, just don't go crazy, where "going crazy" means assigning a set of initial conditions a positive measure where Liouville assigns measure zero. As an instance, since the Liouville measure for any *precise* initial condition is zero, so should the measure for any precise condition be zero in the new measure. This amounts to saying that the ball-dropper, be she right-handed or left-handed or ambidextrous, is not *infinitely precise* in her dropping ability. That is, to say the least, an extremely mild condition.

So even if we replace the Liouville measure over the possible initial conditions with one lumped up to the right, and so break the perfect spatial symmetry of the set-up, still the dynamics and configuration of the board will pick out the precise frequency of .5 as the result of almost all initial conditions. And if we change the configuration of the board by, say, making the pins elliptical and slanting to the right rather than round, then the (infinitely hard) analysis should give us some other, higher limiting frequency of Rs to Ls for almost all initial conditions. It is in this sense that the typical limiting frequency is determined by the microdynamics and the configuration of the board rather than by details of the ball-dropper. The ball dropper—aside from dropping the ball into the hopper and not being infinitely precise about it—evaporates from the analysis. And

the *beliefs* of the ball-dropper, or of anyone else for that matter, never came into the analysis in the first place.

One might wonder, if any measure absolutely continuous with Liouville measure defines the same class of typical behavior, why the Liouville measure in particular has been adverted to. That is because the measure—being equivariant—plays well with the dynamics. The mathematical analysis of the situation using any equivariant measure will be much simpler and more straightforward than using any other.

In the last two paragraphs we introduced the term “typical” in the phrase “typical limiting frequency”. In this setting, “typical” has a precise mathematical meaning: a *behavior* is “typical” if the set of initial conditions yielding that behavior has measure 1 in Liouville or any other absolutely continuous measure. In this precise sense, the set of initial conditions yielding that behavior is “almost all” or “overwhelmingly most” of the initial conditions consistent with the description of the chance set-up. Of course, as Sheldon Goldstein often insisted to me, the “bad” initial conditions (i.e. initial conditions leading to some other frequency or to no limiting frequency at all) are still there. They exist, and in that sense a ball thrown in the hopper *could* (as a matter of physical possibility) do anything. No amount of talking or of mathematical analysis will make the “bad set” of initial conditions disappear entirely. Nonetheless, if the set leading to limiting frequency .5 is almost all of the possible initial conditions, that is an extremely important result. It is a result that creates as much scientific explanation and understanding as one could ever hope for in such a situation. That fact means both that one should *rationally expect* a ball thrown into the hopper to display a limiting frequency of .5, and that one should not feel obliged to *seek any further explanation or question one’s physical theory of the situation* if such a frequency is observed. Furthermore, one should *feel puzzled* and *seek some unknown physical interference* if any other frequency is observed. If one were presented with an infinite quincunx and threw a ball in it, and the ball displayed any other behavior than a limiting frequency of .5 Rs, then one would rightly reject the claim that that quincunx is “fair” and the pins exactly cylindrical, etc., even though the outcome is *consistent* with it being fair. And if the board *is* fair, then all we can say is that you have suffered a terrible case of epistemic bad luck. Drawing the only rationally acceptable conclusion from the available evidence, you have been led astray. Such rational bad luck is an unavoidable, if distressing, possibility in life. For example, if two students turn in word-for-word identical essays, one concludes that at least one of them has cheated, even though it clearly *could*

just be pure coincidence. But, as we say, “what are the chances?”, and we rationally conclude cheating has occurred. If it hasn’t, epistemic bad luck for us (and worse luck for the students). That’s just the nature of our epistemic situation, and we should accept it.

So in this account, rational credence comes in only at the *end*, as a consequence of the physical analysis. It plays no role in the analysis itself. The analysis yields the conclusion that a certain observable behavior is typical for a generic chance set-up, and comes to that conclusion using only resources from categories A and B above, not C.

As a conceptual matter, it is important to note that what is “typical” in this account is a *behavior*, and in particular a *generic behavior* (i.e. a behavior displayed by systems with different precise physical states). It is never accurate to say that any particular initial condition is itself “typical”. Typical behaviors are displayed by a set of initial conditions with measure 1, and no particular initial condition has measure 1. Every particular initial condition will display many typical behaviors and also some atypical behaviors. For example, every initial condition in our quincunx yields a *specific* sequence of Rs and Ls. And displaying that particular sequence will be atypical: it is only displayed by a set of measure 0.

At this point, the most conceptually difficult issue confronting this approach becomes evident. Suppose (contrary to real possibility) one were presented with an infinite quincunx and threw a ball in and could somehow have access to the precise infinite sequence of Rs and Ls that results. I have claimed above that if that sequence fails to have limiting frequency of .5, then one ought to conclude that the board is not, in fact, fair (in the sense of having a perfectly symmetric configuration of pins). And that is because limiting frequency .5 is typical for such a configuration and any other frequency (or no limiting frequency at all) is atypical. So far so good, and everyone should agree. But now the objection is raised that the *particular* sequence that occurs (no matter what it is!) will be atypical, and so one might predict that *that particular* sequence would not occur on the same grounds. Why, then, should *that* not be taken as rationally compelling evidence that the board is not fair? In short, in some cases displaying behavior that would be atypical for a system of a certain sort is evidence against the hypothesis that it is of that sort, while in other cases it is not. How does one separate the sheep from the goats here?

This is an excellent question, and one to which no complete precise answer exists (to my knowledge). Here are some observations about it.

First observation: of course one cannot cast doubt on the claim that the quincunx is fair merely because the exact sequence is atypical, because that is true of *every* exact sequence, so by that principle one would have grounds for doubt *no matter what* the outcome! That is clearly bad methodology. Nothing similar happens with the frequency: if the board is fair then one expects the limiting frequency of .5, and that result confirms the hypothesis. An evidential game that the system can't win is a rigged game.

Second observation: the relevant description of the frequency behavior—"limiting frequency .5" is quite compact. The description of the exact sequence of Rs and Ls is infinitely long and typically not finitely specifiable in any way.

Third observation: there is no alternative hypothesis under consideration before the ball is released that would make that particular sequence typical. If there were, for some reason, then of course the obtaining of that particular sequence would massively confirm the alternative and reduce the credibility of the hypothesis that the board is fair.

All of these seem relevant to the fact that we would, indeed, take the occurrence of an atypical frequency (or lack of limiting frequency) as strong evidence against the hypothesis that the board is fair, but not take the occurrence of the particular atypical sequence as such evidence. These observations do not, of course, amount to a complete theory of how we make such distinctions, but merely point in the direction of how such a theory might be pursued.

So the situation with respect to the infinite quincunx is reasonably clear, and in particular the mathematical aspect of it is perfectly sharp due to the appeal to sets of measure 1 and 0. But that mathematical sharpness depends critically on the board being infinite. For any finite board, every sequence of Rs and Ls will be finite, and so have a particular frequency, and the set of initial conditions leading to each possible frequency will have a positive (Liouville) measure. The definition of "typical" as "produced by a set of initial conditions of measure 1" will therefore yield no useful results at all. And the appeal to absolute continuity with respect to Liouville measure cannot be used to extend the result from Liouville to an infinite class of measures. Nonetheless the infinite case gives us the correct logical skeleton to appeal to. What we need to understand is how to retain the basic anatomy of that skeleton while de-idealizing from the unrealistic infinite to the realistic finite case.

De-Idealizing the Quincunx

No actual quincunx—and no actual outcome data from any experiment—is infinite. So in order to make contact with actual scientific practice, the case discussed above must be de-idealized. The de-idealization introduces both a degree of arbitrariness into the definitions and analysis, and also requires the use of uncertainty bars or intervals where point values appear in the infinite case. Both of these adjustments are important conceptually, and also bring the analysis in line with actual scientific practice. As we will see in the next section, these new aspects of the account also appear when the general question of practical reasoning arises.

Consider a finite quincunx of, say, 1,000 rows. In such a case, clearly no behavior of the marble through the device will be “typical” in the sense of “produced by a set of initial conditions of measure 1”. Every possible sequence of Rs and Ls, and so every possible frequency of the form $N/1,000$, will have some non-zero chance of occurring. So if the whole conceptual apparatus is to have any non-trivial application at all, we need a less demanding definition of “typical” behavior.

Of course, we know how that goes. Instead of probability 1 (in the appropriate measure), we have to choose some threshold close to 1. How close? There are dangers in both directions. If we chose a threshold too high, then we run into the same problem as keeping it at 1: no non-trivial behavior will be either typical or atypical so no conclusion from data can ever be drawn and no rational expectations formed. Also, the higher we set the threshold, the more data that needs to be collected to get enough to be significant, and that entails practical difficulties. But if we set the level too low—at .95, to take a salient threshold—we will get too many evidential false positives and false negatives. We will give too much weight to crackpot theories that just happen—by sheer luck—to make an accurate prediction.

In high-energy physics, a high but still manageable threshold—5 sigma or .99994—has been set for data to count as proving the existence of a particle. That threshold is too high for many other sciences, which cannot generate and analyze such large data sets. But let’s take that as our test case—the adjustments required by a lower threshold are just a matter of pure math.

So: let us say in that case of our 1,000 row quincunx that some behavior is “typical” if it is produced by .99994 of the possible initial conditions, as measured by Liouville measure. Certainly, 99,994 out of 100,000 counts as “overwhelmingly most” by everyday standards. We are advised to expect typical behavior (subject to the sorts of considerations mentioned above) and to consider

the failure of a system to behave typically according to some theory as grounds to revisit or doubt the theory. But what, in the de-idealized case, will count as typical?

Clearly, no *precise* proportion of Rs to Ls in the sequence of a trajectory will be typical. However, certain *ranges* of frequencies will be. And since more precise predictions are preferable methodologically to less precise ones, we ask after the *narrowest* range of frequencies that the hypothesis renders typical. Due to the symmetries of the quincunx (if it is fair), the narrowest such range will be centered on .5 and will extend out in both directions as far as it must to catch .99994 of all the possible trajectories. For our 1,000-row board, that yields a range of .435 to .565. According to the chosen 5 sigma standard, it is typical for the frequency to fall in that range. Any result outside that range provides good grounds to doubt that the board really is fair.

Raising the standard for what counts as “typical” broadens the size of the relevant interval, and therefore makes the prediction of the theory less precise and useful. Adding more rows to the board—or sending more than one marble down the board—narrows the range of the interval of typical behavior. The details of how this works and how to handle the epsilonics are the stuff of standard statistics, and on a solid mathematical footing. Our only question is how it bears on the conceptual issues surrounding “deterministic chance” and scientific method.

What we have just seen is that although fundamental chances—chances directly part of the fundamental dynamics—are point-valued, typical frequencies as they arise in generic chances are (in realistic cases) interval-valued. In a collapse dynamics for quantum theory, for example, the proposed law will assign an exact value as the chance of a collapse in unit time. But the generic chances in real cases do not trade in point values. The notion of “typicality” only attaches to intervals of values, and size of the interval itself depends on an arbitrary choice of threshold. And the choice of a reasonable threshold is also somewhat vague, with many different choices being defensible. These differences in kind between fundamental and generic chance follow from the nature of the chances themselves.

The vagueness about the threshold should not bother us any more than the pervasive vagueness of almost all concepts does. The term “shirt” is obviously vague: at what point does a thinner and more ripped piece of fabric draped over the body cease to be a shirt and become a mere rag? There is no sharp answer, but restaurants with the rule “no shoes, no shirt, no service” are able to function perfectly well despite that.

The final consequence of de-idealization arises because we can no longer say that whatever results follow for Liouville measure over the space of initial conditions also follow for any absolutely continuous measure. It is still the case that the Liouville measure will be the preferred measure to *carry out* the analysis with because it is equivariant—it plays well with the dynamics. But since the Liouville measure may not be a particularly realistic reflection of the chance set-up (as with our right-handed ball dropper), we also want to be able to relax the dependence on that particular measure for defining “almost all” or “overwhelmingly most” as much as possible. And here a brand new relevant aspect of the dynamics comes into play: chaos.

Intuitively, a chaotic dynamics displays a very sensitive dependence of trajectories to small changes in initial conditions. That is, in a non-chaotic dynamics (such as the two-body gravitational system) a small perturbation in initial conditions (changing the initial position or momentum of a planet just a bit) typically yields a slightly perturbed evolution (the later positions and momenta of the planets will only be slightly different). In a chaotic dynamics (such as hard spheres colliding elastically in a box) a slight change in initial conditions usually results in a huge change in the later configuration of the system because the differences are amplified by the collisions. As has been often pointed out, if the dynamics of a system are chaotic, then relatively smooth changes in a probability measure over the space of initial states will yield only tiny changes to the measure over generic outcome states. Let’s see this in action in our quincunx.

We have characterized the outcome of sending a marble down the board at three different levels of description. One is the completely exactly physical description given by the fundamental microdynamics. Another is the sequence of Rs and Ls instantiated by that precise trajectory. And even more generic is just the proportion of Rs, or limiting frequency (if any) of Rs of the trajectory. Either of the last two is a generic description of the outcome that can be realized by infinitely many distinct precise trajectories.

For each such generic outcome, there will be a set of precise initial conditions that lead to it. The measure of that set reflects, in some sense, “how likely” it is given the chance set-up. But beside the measure of that set, there is also a question of how the set is *distributed* in the space of initial conditions. It is easiest to grasp the situation by visualization. Suppose that the initial condition space is two dimensional: every marble starts out in some particular location above the hopper and with an initial momentum near zero (the marble is “dropped”, not “thrown”). In the two-dimensional space of all initial conditions, paint the “good” set—the set that lead to the generic

outcome in question—green and the “bad” set red. To say that “overwhelmingly most” of the initial conditions (as measured by, say, Liouville measure) result in a particular outcome is to make a remark about *how much* of the space is green and how much is red. But to say that the dynamics is relevantly chaotic is to say something else: it is to say that the red and green points are thoroughly intermixed at many scales. It is not that all the red, for example, is bunched up in one corner—no matter how small—of the space. Furthermore, the proportion of red area to green area can be fairly constant under coarse graining. That is, take a sphere or cube (or other compact shape containing a high volume for a given surface area) and plunk it down anywhere in the initial condition space. The red and green are nearly uniformly distributed in a coarse-grained sense (where the degree of coarse graining is measured by the volume of the sphere or cube) if the proportion of red to green is about the same no matter where the sphere or cube is centered or how it is oriented. If all the red were bunched up in one corner it would not be uniformly distributed in this sense. A chaotic dynamics can yield such a uniform coarse-grained distribution.

Now: if the distribution of initial conditions leading to a particular generic outcome is uniform in this coarse grained sense, then we can make an assertion in the de-idealized case analogous to that in the idealized infinite case. In the idealized case, we could say that any behavior that is typical with respect to Liouville measure is also typical with respect to any absolutely continuous measure. In the de-idealized case with chaos and a coarse-grained uniform distribution we can say that any behavior that is typical with respect to Liouville measure is typical (or nearly so) with respect to any measure that differs only *slowly* from Liouville. That is to say that the difference between Liouville and the new measure changes only slightly over the distance scale introduced by the coarse-graining. Or, more technically, the first derivative of the difference between the measures is everywhere small over the coarse-graining scale. In layman’s terms, although the new measure may differ in different regions substantially from Liouville, the spatial changes in it are mild: it does not go up and down like a roller-coaster over the coarse-graining scale.

In the infinite case we could say that if you don’t like Liouville measure as a representation of the possible outputs of the chance set-up, fine: change it but don’t go crazy by putting positive measure on a set of Liouville measure zero. In the finite case we say: if you don’t like Liouville measure as a representation of the possible outputs of the chance set-up, fine: change it but don’t go crazy by having the new measure change dramatically over the scale of the finest coarse

graining that keeps the proportion of green to red uniform. As long as the change in the measure satisfies this condition, the result for what counts as “typical” behavior will not change either. So that judgment is not hostage to the precise choice of an equivariant measure, even though that will be the most convenient to use for the analysis.

This sort of analysis has a long history going back at least to the work of Henri Poincaré in 1896.²⁰ It has been explicated in many places in more detail than is required here. The main conceptual point I want to insist on is that the chaotic dynamics, if it obtains, 1) is a perfectly objective physical fact about the system; 2) is only one of several other perfectly objective facts that may be adverted to in defining generic chances; and 3) makes the de-idealized case even more similar to the idealized case which has sharp mathematical results that are independent of any arbitrary choice of thresholds.

At the risk of being tedious, let me repeat the upshot one last time. Generic chances are chances ascribed to generically specified events in generically described circumstances. Because the relevant conditions are not specific, the microdynamics cannot be immediately brought to bear on the problem. The filling in from the generic level of description to the specific must be done via a measure over the space of initial conditions, and our main conceptual issue is what the nature and origin of that measure is, and in particular whether it must have some *subjective* origin in the *credences* of some agent. I have argued that this is not the case for several different reasons. First, there are measures that are picked out by relation to the objective dynamics by being equivariant, and that has nothing to do with anyone’s credences. Second, there are measures that are derived from spatial and temporal measures already postulated by the physics, and that has nothing to do with anyone’s credences. And third, given the use that these measures are put to—in defining a notion of “typical” behavior—the concepts are not even restricted to either of these sorts of objective measures: one would get effectively the same results using any of an infinite class of measures which are, in a well-defined way, similar to these (e.g., in the infinite case absolutely continuous with them). This last puts the burden on the other foot: rather than the proponent having to *defend* appeal to typicality as defined here, the opponent would have to put forward positive grounds to prefer some other measure that falls outside of this quite widely extensive class.

Finally (and relevant to the next section), in the realistic finite case what counts as “typical” in terms of frequencies in the observed data is necessarily a *range* of outcomes rather than a point

²⁰ See von Plato 1983 for the history.

value. That circumstance has important ramifications for the question of what one should *expect* when confronted with a choice of various different courses of action and has to make the practical decision of which to pursue.

Practical Reason

Our final topic is practical reason. Practical reason is so deeply intertwined with both rational credence and objective chance that it is essentially impossible to discuss the latter without invoking the former, and we have done so from the outset. In order to illustrate that the chance of the dart hitting either p or q is strictly greater than the chance of it hitting p we offered the agent a choice between a ticket that pays only if p is hit and two tickets: the one that pays if p is hit and the other if q is hit. Not having a definite preference for the latter over the former is simply insane. There would be no way to defend it. (Again, we are just throwing darts at random: no Newcombe-like funny-business where the choice is statistically correlated with the throw.)

This case is so clear and undeniable because it follows from a clear and undeniable principle: Dominance. Dominance goes like this: Start with set of outcomes that could eventuate. Each outcome leads to a payoff if one takes Option A and a payoff if one takes Option B. Order rank the payoffs in terms of preference. (We assume that payoffs can be strictly order-ranked: if there are incomparable payoffs, then that obviously leads to complications but they have nothing to do with chance or credence.) Now: if for each possible outcome, the value of the payoff is at least as high—and sometimes higher—on option A than on option B, then A should be chosen. Period. End of story.

Dominance does not even mention probability, it merely says that no matter what happens one is better off having chosen A than B. So choose A.

One can extend this principle in an unobjectional way to cover a wider range of cases where one has available judgments of relative probability or chance. If for every possible value of a payoff, Option A provides at least as high a probability for a payoff at least that good—and sometimes a higher probability—as Option B, then Option A should be chosen.²¹ Call this

²¹ Thanks to Adam Elga for this formulation.

principle “Stochastic Dominance”. Stochastic Dominance only requires for its definition that the values of the payoffs and the objective chances be order-ranked, not that they have ratios defined between them. So it can also be implemented in a setting where Kolmogorovian probability measures have not been introduced.

The judgment about the likelihood of an outcome given the choice of an action is exactly a judgment about a future subjunctive conditional: If I should do A, O would happen. The degree of rational credence in that future subjunctive is just the degree of credence in O conditional on the supposition of A. So whether one believes that A dominates B is straightforwardly determined by one’s degrees of credence in those subjunctive conditionals.

Dominance and Stochastic Dominance unambiguously resolve many cases that are troublesome on the “expected value” approach due to “probability zero” events (according to some Kolmogorovian measure). Sometimes Dominance or Stochastic Dominance, like the Euclidean Principle, resolves the issue directly.

But just like the Euclidean Principle, Dominance and Stochastic Dominance are only rarely applicable. If all decisions could be settled by these principles, then decision theory would be trivial. So the next question is what to do when neither option dominates the other.

The usual answer here is to appeal to expected utility: Don’t just order rank the possible outcomes and probabilities, but rank them on a ratio scale. Next, weight each outcome by the credence one assigns to it (which will be the objective chance one takes it to have, if one has an opinion on that), calculate the “expected value”, and then choose the option with the highest expected value. One could basically keep to the same scheme using relative credences as we have explicated them. Rather than multiplication one would use ratios of credences, with the aim of order-ranking the “expected values”. There are a few technical details to be worked out, but nothing really hard.

However, all of that is wasted effort, because the advice to choose the option with the highest “expected value” was bad advice to begin with. So although one could implement this scheme to produce a theory of rational choice one shouldn’t.

The problem is familiar but has not been taken seriously enough. The classic “paradox” illustrating the problem is the St. Petersburg paradox, so let’s start there. A denumerable infinity of fair coins will be flipped (we do them all at once to avoid issues about waiting for the payoff). The coins are not merely denumerable, they have been enumerated. The payoff structure for a

ticket is this: Let N be the number of the lowest-numbered coin to come heads. That yields a payoff of $\$2^N$. In the event that none comes heads, let the payoff be whatever you like: in the end it is irrelevant.

What would be a fair price to pay for that ticket? According to the usual approach, that would be the “expected value” of the gamble but the expected value, as standardly calculated, “is infinite” (i.e. the calculation diverges and the expected value grows without bound). That particular feature is actually irrelevant to the puzzle: for any proposed value of a ticket, the calculation could be cut off at a point where the expected value straightforwardly exceeds it, so according to the principle one “ought” to be happy to pay that much. But it would be completely insane to pay, say, \$1 million for such a ticket.

It gets worse. Change the payoff scheme so that unless the first million coins all come tails, the payoff is zero, and if they do then the payoff is $\$2^{(N-1,000,000)}$, with N the number of the first heads. Again, let the payoff for all tails be whatever you want. The expected value of the gamble would still diverge to positive infinity, and one should be willing to pay any amount for a ticket. But it would be complete lunacy to pay *anything*, even if the “expected value” is positive.

The common attempts to defuse this disastrous result involve denying that the value of a payoff scales with its monetary value. Of course, one could then increase the monetary values to compensate so long as there is no roof to the value. But as with the Banach-Tarski paradox, the wrong suspect has been fingered: the issue is with the “expected value” rule itself, not anything with its implementation.

After all, why in the world should one *want* to maximize the “expected value” of the choice? One has no direct interest in the “expected value” of a choice, one has an interest in the *value* of the outcome! In a case of decision-making under uncertainty, one doesn’t know what that value will be, because one doesn’t know what the outcome will be. But that alone does not mean one can’t know how to act: for example, if one option dominates the other then one knows one will not be worse off with it no matter no matter the outcome.

The real trick here is a sort of cognitive illusion created by a clever choice of nomenclature. There is a certain quantity calculated: call it the “credence-weighted-value-sum-of-possible-payoffs” for each choice, because that’s what it is. What it most certainly is *not* is the “expected value of the outcome” in the obvious sense: the value of the payoff one expects to receive! If it *were* the value of the payoff one expects to receive, then it would make at least *prima facie* sense

to opt for the action that has the better one, but there is no even *prima facie* motivation to choose the action with the highest credence-weighted-sum-of-values-of-possible-payoffs.

The point is obvious. Consider a lottery with 10,000,000 tickets, with the tickets costing \$1 each and the payoff \$9,999,999. The “expected value” is $-\$.0000001$, so the rule says not to buy the ticket. Now increase the payoff by \$2 to 10,000,001. The “expected value” is now $+\$.0000001$, so you ought to buy the ticket. That already seems rather strange, but the first point is that neither $+\$.0000001$ nor $-\$.0000001$ is the amount of money you *expect* to get. In fact, you *know* that neither of those is the payoff that will eventuate: not only is this not the value you expect from the option, it is a value you know will *not* result from the option. So the term “expected value” is completely inappropriate. It makes the rule sound innocuous when it isn’t. (After all, it is the rule which recommends almost certainly bankrupting yourself buying a ticket for the St. Petersburg lottery!)

What the “expected value” really is, of course, is the *average* value you would expect to get were you to play the same lottery over and over and over enough times. And indeed, if you were to play these lotteries over and over many, many, many millions of times, then you could reasonably expect in the long run to lose $\$.0000001$ per play in the first lottery and gain $\$.0000001$ in the second. That’s the sort of reasoning that casinos engage in, and because they do take many, many, many bets at favorable odds they regularly and predictably make money. But if you have no intention or no resources or no opportunity to play such a lottery billions or trillions of times, it is hard to see what counterfactuals about what would almost certainly happen if you did play that many times have to do with anything. (If you had that much money, then why not buy up *all* the tickets? In the first lottery you would expect—with certainty!—to lose a dollar and in the second to win a dollar. Then the advice makes perfect sense: do the latter and avoid the former.)

And if you could play the St. Petersburg game enough times (don’t try calculating how many) then you could indeed reasonably expect to make money off it, no matter the price of a ticket. But you won’t, so why care about the counterfactual?

In fact, leaving St. Petersburg aside, it is perfectly clear what you ought to expect if you buy a ticket for either the \$9,999,999 lottery or the \$10,000,001 lottery: you ought to expect to be out a buck with nothing to show for it. You ought to expect that to the *very same degree* in both cases. So a pretty reasonable piece of advice would be: if you like money, don’t play. In the proper sense of the term, the expected value of buying a ticket in both cases is $-\$1$.

Let's introduce a piece of terminology: let's say that you "expect" something to happen when your credence in it is greater than your credence in a fair coin being flipped and coming heads at least once in 20 flips. The official probability of such an event not happening is 1 in 1,048,576, a number so small that no one really has in intuitive sense of what it means. But everyone has played with coins enough to know that while flipping a coin and getting tails 20 times in a row is theoretically possible, as a practical matter it just ain't gonna happen. Such an event not happening is what used to be called, in philosophical circles, "morally certain". That is, the certainty, while not absolute (Top) is enough for practical (moral) purposes.

The choice of 20 coins flips for the standard is obviously somewhat arbitrary and could be adjusted to 15 without harm. But it can't be reduced to 10: ten tails in a row is an eventuality which, while unlikely, is not negligible.

With that terminology in place we can certainly propose this rule: if you expect one action to have at least as good an outcome as another, you are permitted to take it and if you expect it to have a better outcome you must take it. That is a much more wide-ranging piece of advice than either Dominance or Stochastic Dominance. It tells you, for example, not to buy a ticket in either of the 10,000,000-ticket lotteries mentioned above.

That still leaves many, many situations unsettled. When you throw a pair of dice, none of the twelve (or thirty-six) possible outcomes—or any disjunction of them save the disjunction of all of them—is expected. That's why one uses dice in games: to create situations with a degree of uncertainty that circumvents the use of the rule just announced.

But if instead of just one throw of the dice one intends to make many throws, then the situation is different. One can, for example, certainly expect not to throw ten boxcars in a row.

This rule—which is essentially to treat (in a circumscribed way) events that are expected in the sense defined above as if one knows they will happen goes by the name Cournot's Principle. It is an excellent source of good advice. Cournot, for example, would advise against paying *anything* for the second St. Petersburg lottery mentioned above, rather than paying any amount at all as the "expected value" rule commands. The sense in which one treats expected events as if one knows they will happen is circumscribed because the propositions taken as "certain" are not closed under logical deduction. Given a lottery of 10,000,000 tickets, one is morally certain of each ticket that it will lose, and perfectly certain that one or another of them will win. But this failure of closure under deduction is a mark of all rational belief. Every rational person believes

that he or she holds some false beliefs. It would be crazy not to. That only yields the Preface Paradox if one insists that rational belief be closed under deduction, and the only conclusion of the paradox is that you should not insist on any such thing.

So if we adopt the Cournot Principle decision theory to supplement Dominance we expand the range of applicability of the theory considerably. But it still will not cover many one-off decisions, even when one can calculate the “expected value” of the options. This is not a sin of commission, like telling you to pay \$1,000,000 for the St. Petersburg ticket, but a sin of omission. The rule is simply silent in many one-off cases.

Here is something that would help with that. Imagine one will be confronted with the same decision some large number of times. And imagine you have to bind yourself to a rule: you have to make the same decision each time. Then apply Cournot’s principle and if the rule makes recommendations follow them in the one-off case. For the moment, we will set the “large number” of plays at 50,000. For reasons to be given, that is a maximum: it could be reduced but not increased.

To make clear what I mean, let’s consider the two versions of the St. Petersburg lottery above in the one-off case.

For the first lottery, if one buys a ticket one can expect—given our standard of expectation—to win at least \$2. That is so even if the payoff for all tails is \$0. One isn’t absolutely certain, but morally certain to win at least \$2. Therefore, one should be willing to pay at least \$2 for a ticket. On the other side, one is morally certain *not* to win more than $2^{20} = \$1,048,576$. So one should *not* pay more than that. That leaves, of course, a wide avenue of discretion, where the rule makes no recommendation. But since the “expected value” rule recommends paying *anything*, already we are, as it were, infinitely better off! These limits are locked in: no further considerations can alter them.

Of course, intuitively we think we should be willing to pay more than \$2 for a ticket and unwilling to pay anything close to \$1 million. Let’s see if we can justify that intuitive judgment.

Fictively imagining playing the lottery a large—but not infinite!—number of times yields some expectations of the following form. In terms of standard statistics, we want to know what we can predict with about .999999 degree of certainty. Since this is all very rough, we will just do a rough calculation for illustration.

If we play the lottery 50,000 times, the chance of getting 23 tails in a row even once is under our threshold, so we can eliminate any payoffs over \$8,388,608 from consideration. The first big payoff we can't eliminate is \$4,194,306. The chance of getting that once *and* any other very high payoff falls below our threshold, so we can ignore that too. We do expect to get many smaller payoffs, and indeed at 50,000 repetitions, we expect the observed proportion of small payoffs to be very close their objective chance. Since the chance of a payoff of \$2 is .5, we can reasonably expect about 25,000 such outcomes. Similarly, about 12,500 \$4 outcomes, 6,250 \$8 outcomes and so on. Since the number of successes halves while the payoffs double, each of these contributes the same amount—\$50,000—to the cumulative total. This is just as the standard expected value calculation returns, because it is rational to expect about the “expected value” to be the *average* value over enough repetitions. So each of the low-payoff results adds about \$50,000 to what one could reasonably expect, and so about \$1 to one can expect the average to be.

In the standard calculation, this same reasoning applies to *all* possible payoffs, so each adds a dollar to the expected average. And since there are infinitely many possible payoffs, that yields complete and total disaster. But we have already averted that disaster since we disregard every result with more than 22 tails and a payoff of more than \$4,194,306. So if we were to treat *all* the remaining possible outcomes as if the “expected value” were really to be expected, we would have an expectation of a total payoff of \$50,000 for each of 23 non-negligibly possible outcomes. That, of course, yields a total payoff of $23 \cdot \$50,000$, and an *average* payoff per play of \$23. This back-of-the-envelope calculation, coupled with the first, yields a maximum rational amount to pay at \$23, which is not only infinitely better than the standard result, it is intuitively quite reasonable for a maximum. It is however, a bit high. We can reasonably expect to average about the “expected value” if we play 50,000 times for the low-dollar payoffs since there will almost certainly be a large number of low-dollar payoffs. But as the payoffs grow, so do the statistical uncertainties. At 50,000 plays, we can expect *not* to get a payoff of \$8,388,608, but can't *expect* to get a payoff of \$4,194,306. That payoff is neither rationally forbidden nor rationally expected. At the very low end, though, between 24,475 and 25,525 \$2 payoffs is again rationally expected. The calculation this “confidence interval”, the smallest range of outcomes one is sure of getting at a 999,999 out of 1,000,000 standard of certainty, is the key to the best calculation. Let's do the complete St. Petersburg example.

But rather than calculate the probabilities for the \$2 payoffs, we should calculate the rational expectations for the higher payoffs. Since every payoff is at least \$2, we can use a total of \$100,000 as a floor and then figure out the rational expectations for the smallest and largest amounts by which the floor will be exceeded.

If one plays 50,000 times, there is a .999999 chance that the outcome TH will happen between 12,040 and 12,960 times. So one can expect those payoffs to contribute between \$24,080 and \$25,920 extra winnings. The expectation for TTH, with an extra payoff of \$6, is that it should happen between 5900 and 6600 times, for a minimum excess winning of \$35,400 and maximum \$39,600. The TTTH outcome should occur between 2870 and 3380 times, for an extra payoff between \$40,180 and \$47,320. We already see the range of the values spread as the number of expected wins decreases, resulting in a larger range of outcomes that ought not to be surprising.

Continuing in the same vein, the range of wins from TTTTH is rationally foreseeable as 1380 to 1750 for extra winnings \$41,400 to \$52,500; TTTTTH 650 to 920 for \$40,300 to \$57,040; TTTTTTH 290 to 490 for \$36,540 to \$61,740; TTTTTTTH 120 to 270 for \$30,480 to \$69,120; and TTTTTTTTH 50 to 147 for 25,500 to \$74,970. At this point the range of values, due to possible statistical fluctuations, has become so extreme that the calculations are essentially useless. In the later cases, of course, the average of the minimal and maximal excess winnings is about \$50,000, so we can just use that as an expectation for the rest of what we consider possible outcomes.

In a pessimistic mood, one might choose the lower bound in each case we calculated. If we do that, we arrive at a total payoff of $\$100,000 + \$12,040 + \$24,080 + \$35,400 + \$40,180 + \$41,400 + \$40,300 + \$36,540 + \$30,480 + \$25,500 + 10 \times \$50,000 = \$858,920$. That would suggest one might reasonably be willing to pay \$17.18 in a pessimistic mood. The optimist, taking the upper bound, would be amenable to paying \$20.76.

These are, of course, rough calculations about a question for which there is no perfectly accurate correct answer. There just is no uniquely amount a single play of the St. Petersburg lottery is “rationally worth”. But a rough calculation using reasonable numbers yielding the answer \$19 plus or minus a couple of bucks is literally *infinitely* better than the standard approach and gives an intuitively acceptable response. I asked people, some familiar with the lottery and others new to it, how much they would be willing to pay for a single ticket. The responses were: \$3 or \$4, \$4, \$6, \$6, at least \$8, \$10, \$10, \$20, \$20, \$20, \$20, \$35, around \$50, \$127.99 and \$500. The most

common response was \$20, an answer that suggests that the Cournot approach, implemented in this way, is a spherical Earth rather than the “expected value” spherical cow. All of the higher offers were made by people familiar both with the lottery and with how to calculate expected values. I doubt that the uninitiated would ever venture such high offers.

We have made two somewhat arbitrary choices in implementing the Cournot strategy: the “20 tails in a row” standard for something so unlikely as to be practically negligible, and the “50,000 repetitions” standard for the fictitious repetition in a one-off case in order to get more useful advice. The latter of these is clearly on the high side. If a human being does something once a day every day of their life, they end up doing it much less than 50,000 times, and a decision made more than once a day would not be considered a “one-off”. At a greater frequency, such as the hundreds of decisions made in playing poker for an evening, the object of the choice is the adoption of a rule actually to be repeatedly used, and should be treated as such. The chance of getting 20 tails in a row is essentially one in a million, which is the phrase commonly used for something one thinks ought not to be taken seriously as a consideration when making decisions. That standard could be dropped to 15 tails (about 1 in 33,000—once in a lifetime at one chance per day!) or even 10 (about 1 in 1,000) and the fictitious repetitions to 500. These low-end values change the numbers a bit, in two ways. First, the price per ticket falls as the possible higher payoffs are neglected. Second, the range of acceptable offers widens as the imagined repetitions are reduced: there is less certainty and more statistical noise. Using those values, we get an allowable range of offers from \$4.26 (for the pessimist) to \$40.09 for the optimist. That range takes in all but the most extreme of the “intuitive” answers. Quite a respectable conciliation of the theoretical and gut-level judgments about what is rational. And—after all—evolution must have equipped our guts with a pretty reliable sense of what sorts of risks are worth taking.

One more example of a one-off decision. Suppose someone offers a reward of \$10,000 if you survive a single game of Russian Roulette, with a 1/6 chance of dying. Most people would consider agreeing to such a thing to be insanity, even though most of the people who do agree walk away \$10,000 richer, a not insubstantial sum. And indeed, for most people increasing the payoff would be irrelevant: making the potential winning \$100,000 or \$1,000,000 would not change their minds. At 1/6 odds, neither dying nor becoming rich is an expected outcome by Cournot’s standards, so the simplest application of the rule yields no advice at all. That is a bad result.

The fictive repetition, though, does. Indeed, if one sets the number of fictive repetitions to 100, much less 50,000, one would already be morally certain to be killed at least once. “Averaging out” 100 results means taking the badness of death into account, which will overwhelm the monetary gains: the approach advises not to play. Good.

Note that it is essential when considering the fictive repetitions that the different plays be considered statistically independent of each other: for all trials $\text{Ch}(\text{Tails on flip } N | \text{Tails on flip } M) = \text{Ch}(\text{Tails on flip } N)$ if $N \neq M$. The same is true for the reasoning about Russian Roulette. That prevents this sort of reasoning from endorsing anything in the vicinity of Pascal’s Wager.

Pascal, of course, attempted to displace the question of God’s existence from the jurisdiction of theoretical reason—which he deemed incapable of addressing the matter—to that of practical reason. And he there used the expected value to argue for prudential reasons that one ought to believe in God (and a Christian version of God at that) due to the mere possibility of a payoff of infinite value. And while he starts out analogizing the outcome to the result of a coin flip with even odds, he explicitly notes that the value of the infinite payoff overwhelms the probability: if the argument works at all, it works just as well if the theoretical chance of God’s existence is reduced to any finite positive quantity.

Everyone knows that Pascal’s argument is bunk. It has to be: if it established that one ought to act like a Catholic and try to inculcate that theological belief, it would work just as well for Islam or for some version of polytheism, etc. And no one can believe, or even pretend to believe, in them all. Further, the errors Pascal made are multiple. They begin with changing the venue of the dispute from theoretical reason to practical: if we regard the existence of God as highly unlikely (even if not impossible), then we ought not to believe it no matter the possible practical payoffs. The introduction of unbounded payoffs into the expected value scheme is also a red flag: see the St. Petersburg lottery. But even more basic than that: even if one analogizes the situation to a gamble, it is certainly not a gamble that can be repeated *with the results being statistically independent*. If it were, then fictively imagining playing Pascal’s game enough times would yield some outcomes where the Catholic God exists and others where the Muslim God does, with strikingly different practical outcomes depending on the religious practices that had been adopted. Given the metaphysical characteristics of theological fact, the outcomes *can’t* be regarded as statistically independent, so even granting everything else the Cournot approach—unlike expected

value—will never recommend adopting religious practice or belief on Pascalian grounds. Score one more for Cournot.

Indeed, in the case of God's existence, here's a simple result. If one regards the proposition that God exists as less credible than that God does not exist, then one weakly believes that God does not exist. Period. No practical considerations can change that one whit.

Adding no more than 50,000 fictive repetitions to Cournot's advice expands the applicability of the rule considerably but not in the crazy way of the "expected value" rule, which is essentially this rule with an unlimited number of repetitions. In many, many cases, this rule will yield the similar advice as the expected value rule: exactly those cases where the expected value rule gives reasonable advice. But it blocks the obviously unacceptable St. Petersburg recommendation, blocks Pascal's arguments, and tells you not to play Russian Roulette no matter the payoff if you survive. Eminently reasonable.

This seems like an acceptable way to base practical rationality on degrees of credence. Maybe it too produces some paradoxes or obviously bad recommendations, but they are not the usual known ones. That question deserves further investigation.

In Sum

The Euclidean Principle is a non-negotiable requirement for any acceptable theory of rational credence and objective chance, so any theory of either that denies it can be dismissed out of hand. That alone rules out any theory that uses a Kolmogorov "probability measure" over propositions or events or outcomes *in a foundational way*. And switching to a non-Archimedean field, such as the hyperreals, does not solve the problem because in addition to the non-Archimedean features, both rational credence and objective chance should have incomparable pairs of propositions or events or outcomes.

The thing to do is to remove numbers of any kind from the foundations of the account and deal directly with the structure of relative credences or chances, insisting on the Euclidean Principle from the outset. The operations of addition, multiplication and division of numbers used in the standard approaches are replaced with the use of addition-like and multiplication-like operations defined directly on the structure of relative credences or chances, and the use of a ratio

structure also defined directly on it. This yields a fairly powerful theory with the Euclidean Principle built in from the outset. A ratio version of Bayes' Rule can be articulated.

Finally, this can again be enhanced by appeal to Kolmogorov "probability" functions, so long as judgements of zero probability are regarded as having no significance. In decision theory, it can be enhanced by the use of Cournot's principle, which yields a practical decision theory much more acceptable in its recommendations than the standard one.

Start with Euclid, which is undeniable but rather narrow, and expand out from there. Some plausible sounding principles, such as the Humean Principle and some allied symmetry principles must be sacrificed, but the necessity for that sacrifice has been clear for centuries. The only other option is to insist on strict finitism, in which case the conflict does not arise. By "strict finitism" I mean the view that it is metaphysically impossible for any physical item to have infinitely many distinct proper parts. That would rule out both the infinite divisibility of space or time or space-time and the infinite extent of these as well. A slightly weaker view would reject merely the actual existence of any physical item with infinitely many distinct proper parts. Stronger than that would be to deny the physical possibility of such a thing. Strict finitism denies the metaphysical possibility. If physical and metaphysical possibility coincide, then the last two collapse into the same view. Strict finitism—and its weaker cousins—are options that are worthy of closer consideration....but this essay is too long already.²²

²² Thanks to Shelly Goldstein, Adam Elga, and members of the Rotman Philosophy of Physics reading group (especially Nicholas DiBella) for comments and feedback.

Appendix 1: Non-Archimedean and Incomparable Structures in Euclidean Geometry

It may come as no surprise—or at least as not terribly controversial—that credences may be incomparable: in some cases, there just is no fact about which of a pair is stronger. And it may not seem shocking that credences fail to have an Archimedean structure for some subject matters. But on first glance, both of these properties may seem foreign to standard Euclidean geometry. After all, Cartesian co-ordinates can be used in a natural way to create coordinates over the plane, and those coordinates are just ordered pairs of real numbers. The real numbers, of course, have an Archimedean ratio structure and all the one-dimensional number fields, including the hyperreals and surreals, are total orders. Every number is less than, equal to, or greater than every other. It may come as a slight surprise, then, that there are simple and familiar structures in the Euclidean plane that illustrate both non-Archimedean ratios and incomparability. One such structure is the angle.

Any two continuous curves that have only a single point in common, which is an endpoint of both, may be said to form an angle where they meet. If the two curves happen to be straight lines, then it is a rectilinear angle, and rectilinear angles stand in familiar ratio relations to each other. Indeed, every such angle can be represented by a real number (radians, say, or grads) and the sizes of the angles relate to each other just as their representatives do. But in addition to the rectilinear angles there are others. For example, there is the angle formed where (part of) the circumference of a circle meets a tangent. Euclid was aware of these and called them *hornlike* angles. He was also aware that they are non-Archimedean with respect to the rectilinear angles: the angle with which a circumference meets a tangent is “infinitely sharper” than any rectilinear angle. If one made the mistake of trying to define the magnitude of a hornlike angle as the magnitude of the rectilinear angle formed by the *tangents* of the curves where they meet, then one would get the absurdity that the magnitude of a hornlike angle is zero. That is particularly absurd because pairs of hornlike angles can be straightforwardly compared in magnitude. Consider two circles tangent to a straight line at the same point, but with different radii. Let the centers of the circles lie on the same side of the line. The angle formed by the line and the larger circle is clearly “sharper” than the angle formed by the line and the smaller circle: the circumference of the larger lies strictly between the line and the circumference of the smaller, and if we let the radii of the two circles sweep toward the tangent point in synch, the slope of the tangent of the larger circle is always less than that of

the smaller. Indeed, a natural suggestion is to measure the sharpness of the hornlike angles by their radii: the longer the radius the sharper the angle. Thus every hornlike angle is “infinitely sharper” than every rectilinear angle, and the sharpness of the hornlike angles themselves grow without bound.

This behavior was noted by Hume in the *Enquiry*:

Nothing can be more convincing and satisfactory than all the conclusions concerning the properties of circles and triangles; and yet, when these are once received, how can we deny, that the angle of contact between a circle and its tangent is infinitely less than any rectilinear angle, that as you may encrease the diameter of the circle *in infinitum*, this angle of contact becomes still less, even *in infinitum*, and that the angle of contact between other curves and their tangents may be infinitely less than those between any circle and its tangent, and so on, *in infinitum*?²³

Hume goes on to describe these results as “big with contradiction and absurdity”, but of course they are no such thing. They are merely unfamiliar and surprising. They certainly demonstrate that non-Archimedean structures cannot be avoided if one accepts the coherence of Euclidean geometry.

But once one has grasped this example, cases of incomparable angles are not hard to come by. Think of a Euclidean plane fitted out with Cartesian coordinates in the usual way. Let the x -axis be our straight line and consider two curves that meet it only at the origin: $f(x) = x + x^2 \sin(1/x)$ and $g(x) = x + x^2 \cos(1/x)$ for $0 < x < 1$. Let the value of these functions to be stipulated to be 0 at $x = 0$. (They are obviously continuous functions in the normal sense). The graph of each function forms an angle with the x -axis at the origin according to our definition, but there is no way to compare these angles to determine “which is greater”. As the argument approaches 0, the functions oscillate faster and faster around $y = x$, trading back and forth about which is greater than which. So there is no criterion by which one angle can be deemed greater than the other. Both of their tangents, of course, limit to a slope of 1.

²³ Section XII, Part 2.

Further, the angle between $f(x)$ and the x -axis can be added to the angle between $g(x)$ and the x -axis in a straightforward way: their sum would be the angle between $f(x)$ and $-g(x)$

Indeed, even though $f(x)$ and $g(x)$ each forms an angle with the x -axis, according to our definition, and even though they both contain the origin and meet there, they do not form any angle with each other because no pair of segments from them intersect only at the origin. But even though the two curves have no angle between them, their angles with the x -axis can still be added, as noted above. Despite the unfamiliarity and surprising properties of these sorts of curvilinear angles, they can sometimes be subject to the sorts of “additive” operations we have discussed.

This rather simple example illustrates that even Euclidean geometry is shot through with both non-Archimedean structures and incomparable magnitudes of the same generic type. It should then not be much of a surprise to discover both of these characteristics when discussing chances of events where the events themselves are described using geometrical terms, such as the chance that a randomly chosen point be in a specified set of points.

Appendix 2: The Sleeping Beauty Problem

There is a puzzle concerning credence and, at least incidentally, practical reasoning that has gotten much attention since 1990: the Sleeping Beauty problem. It has occasioned quite a large literature and has been understood to present a variety of different problems. Perhaps the most common account of the problem asserts that it raises questions about how to properly update one’s credences in *de se* (or generally indexical) rather than *de re* propositions. Since *de se* propositions are commonly represented by (or identified with) sets of centered possible worlds rather than sets of possible worlds, the problem has been connected with that of self-locating uncertainty. And from there, links have been made to understanding epistemic uncertainty in the face of personal fission, which in turn connects to the problem of understanding probability in the Many Worlds interpretation of quantum mechanics. So the tentacles of the Sleeping Beauty problem stretch over a vast array of topics. But I will argue that the problem can be straightforwardly resolved by careful attention to exactly what questions are being asked and about what assumptions concerning

practical reasoning employed. I will also argue that the problem has nothing in particular to do with *de se* beliefs, centered worlds, or the Many Worlds interpretation of quantum mechanics.

The set-up is fairly simple. Sleeping Beauty will be put to sleep on Sunday night. After she is asleep, a fair coin will be fairly flipped. If the coin comes heads, Sleeping Beauty will be awakened on Monday without being told what day it is, then put back asleep until Wednesday, when she will be awoken and informed of the day. If the coin comes tails, she will be awoken on Monday without being told what day it is, then given a drug to erase her memories of Monday and awoken again on Tuesday without being told what day it is, put asleep yet again and woken on Wednesday just as in the first scenario. Nothing in the circumstances of the Monday or Tuesday awakenings will provide her with any information about which day it is or how the coin fell. Sleeping Beauty is made aware of the entire protocol and can reflect on it before she goes to sleep.

Question: when she is awakened, what should her credence be that the coin came heads? And what should her credence be that it is Monday rather than Tuesday?

There are two main positions that have been argued, called the *thirder* and *halfer* views. The thirder argues that when she awakes, her credence that the coin came heads should be $1/3$. The halfer argues that it should be $1/2$. Implicit in the argument, of course, is some idea of what the “should” in the previous sentences adverts to.

There are many different ways to approach the problem, and different lines of argument, but let’s just sketch a simple halfer argument and a simple thirder argument.

Halfer: Clearly, when she goes to sleep on Sunday night Sleeping Beauty should have credence $1/2$ that the coin—which is yet to be flipped—will come heads. It is a fair coin and will be fairly flipped. But whenever she awakes, she will (by stipulation) acquire no evidence or information about whether it did come heads. Therefore, she should maintain her original credence in the proposition that the coin lands heads: $1/2$. Furthermore, if the coins comes tails, and she is awoken twice, she has no information about which awakening any particular one is. So her conditional credence on it being Monday supposing the coin came tails should equal her conditional credence that it is Tuesday. So when she awakens, she knows that it is one of three scenarios: Monday/heads, Monday/tails or Tuesday/tails. These are mutually exclusive and jointly exhaustive. So her credences should be: $Cr(\text{Monday/heads}) = 1/2$, $Cr(\text{Monday/tails}) = 1/4$, $Cr(\text{Tuesday/tails}) = 1/4$. Her credence that the coin lands heads should be $1/2$ and that it is Monday should be $3/4$.

Thirder: Before she goes to sleep on Sunday, Sleeping beauty knows that there are three possible awakenings that she might later undergo: Monday/heads, Monday/tails and Tuesday/tails. On Sunday, her credence that she will, in fact, undergo each of these is the same: $\frac{1}{2}$. (There is no problem about the three credences summing to more than 1 since the events are not mutually exclusive.) When she awakens, she knows that one of the three events is occurring, but receives no information about which one it is. Therefore, she has no grounds to change her view that all three are equally probable. But regarding a particular awakening, they *are* mutually exclusive and jointly exhaustive: each awakening will be exactly one of the three. So she should regard them as equally likely, and accord them the subjective credence $\frac{1}{3}$. Upon being awakened, her credence that the coin came heads should be $\frac{1}{3}$, and that is it Monday $\frac{2}{3}$.

It appears that the halfer and the thirder irreconcilably disagree. Of course, they both have in mind some sense of “should believe” that is supposed to be supported by their argument: each regards the argument as *rationally compelling*. And that raises the question of exactly what the relevant rationality principles are. But before even confronting a question like that, let’s ask a much simpler and indeed uncontroversial question.

Suppose we ask Sleeping Beauty not for her credence in the proposition that the coin came heads, but for what she would offer as fair odds on a bet that the coin came heads. Just to make it clear, let’s review one way of stating odds and what it means to regard them as fair.

Let’s suppose that in a bet on heads vs. tails, each player must supply a *stake*: the amount of money that player puts in the pot. If the coin comes the way the player calls, then she gets the whole pot: her contribution and the contribution made by the opposing player. The odds are then stated as a ratio H:T between the two stakes. For convenience, we can normalize the H stake to be the value 1, so the ratio representing the bet will be given as 1:T. To regard these odds as *fair*, a player must be willing to take either side of the bet at those odds.

So, for example, for a single toss of a fair coin the fair odds are 1:1. Each player puts the same amount of money into the pot, and winner takes all. What if a coin were loaded or biased to come heads only one time in three on average? Then clearly the heads bettor should not have to put in the same stakes as the tails bettor. In fact, the fair odds should be 1:2. Why? Because in the long term, the coin will come heads half as many times as it comes tails. So in 300 plays at a dollar (on the heads side) a play, heads will stake \$300 and win about 100 times, taking a \$3 pot each

time, i.e. heads will (approximately) break even, and the same for tails. So there is no advantage to betting the heads side over of the tails side or vice-versa.

Good: now let Sleeping Beauty reason about her situation in exactly this way. If asked what she should *expect* from any single bet (not the expected value, but what she should expect), she will say that she can't rationally expect any outcome. But, as we discussed above, she can fictively imagine playing the same game a reasonable number of times and ask what she could expect then. Suppose, for example, they repeat the experiment 1,000 times, and each time she is awakened she has to propose odds for a fair bet. Since she gets no information from the awakening, the odds will always be the same. What should they be?

Well, if the experiment is done 1,000 times and a fair coin is fairly tossed each time, then (about) 500 times it will fall heads and she will be only awoken on Monday, and (about) 500 times it will fall tails and she will be awoken on both Monday and Tuesday. That makes a total of 1,500 bets. If someone bets heads every time for \$1, she will put a total of \$1,500 into the pot and will win (about) 500 times. In order to break even, then, each pot must yield \$3. So the fair odds—the odds at which neither side can expect to make a profit—must be 1:2. These are, of course, the same odds as for a single toss of a biased coin.

Sleeping Beauty can reason this out and see that if she offers any odds other than 1:2 then the opponent, by taking the appropriate side of the bet, can reasonably expect to win money and she can expect to lose it. In that sense, these are the only reasonable odds for her to offer.

So: do we have a solution to the problem, and if so is it a halfer or a thirder solution?

It might at first seem tempting to call it a thirder solution, since the odds Sleeping Beauty deems fair are 1:2, the same odds for a single flip of a coin biased to come up heads 1/3 of the time. But note: although those are the odds she deems fair, the *calculation* that arrives at those odds assumes that the coin is fair and will come up heads (about) half the time. Her odds are *based* on that assumption. So one might call it a halfer solution.

At this point, I think it is unclear what the debate is even about. If everyone agrees that this is the right way for her to set fair odds, then the practical part is done: she knows how to act. Or if someone offers her odds she knows how to act: if the odds are 1:2, take either side indifferently, if they are anything else take the favorable side. An observer unaware of her reasoning might be fooled into thinking that she thinks the coin is biased, but she doesn't. As we have seen, she sets

the odds this way exactly because she *doesn't* think it is biased and expects it to come heads about half the time.

Now: if one thought that she were setting the fair odds by calculating an *expected value*, using her subjective credence for the coin flip to weight the possible outcomes, then one would have to conclude her subjective credence is 1/3. But that just isn't what she is doing. Rather, she is forming an expectation for what the outcome of a longish-but-not-too-long sequence of gambles at various odds would be and setting the odds so she can expect to (approximately) break even: there is no advantage to betting one side over the other. Following this method for approaching practical decision making we get to the fair odds in a different way than is usually done, and the whole problem evaporates. Sleeping Beauty thinks that the coin is fair, with equal objective chance to come heads and tails, and that *therefore* the fair odds for a bet—in this experimental condition—are 1:2.

What is unusual about the experimental condition, of course, is that there is a long sequence of bets but a fresh coin or fresh flip are not used for each bet. Whenever the coin comes tails, there are two bets made on the outcome of that flip, and whenever it comes heads only one bet. So if she were to offer the odds 1:1 she would be taken to the cleaners by an opponent who always bets tails. If the opponent bets the same amount—\$1 for example—each time, he would end up betting twice as much money on the flips that come tails as on the flips that come heads. Hence *because* heads and tails come up (about) the same number of times, he makes money. His losses on heads are more than outweighed by his gains on tails.

Perhaps it is easier to see this way. Suppose a fair coin will be fairly flipped many times, and the odds offered on each flip are the usual fair odds of 1:1. And suppose that the opponent has to lock in a choice: either he has to bet tails for every flip or bet heads. And suppose he locks in on tails. But suppose he has the choice, on each flip, to bet either \$1 or \$2. And suppose beforehand he *knows* what the outcome will be. Then he can make a fortune betting \$1 on the flips that come heads and \$2 on the flips that come tails. *Because* the coin is fair and will come heads (about) half the time, we can easily calculate how much money he will make on average: about a dollar for every 2 flips. The situation with Sleeping Beauty is just the same, except that the opponent needs no precognition. As things are set up, if a bet is made on every awakening, twice as much will be wagered on the tails flips as on the heads flips. So if the odds of 1:1 were offered on every flip, the bettor taking tails would clean up.

Once one sees that this is what is going on, it also become clear that the whole puzzle has nothing at all to do with indexicals or centered worlds or *de se* beliefs, or conditionalization. Here's a simple way to see that.

As before, let the experiment be run 1,000 times with (about) 1,500 awakenings. But instead of Sleeping Beauty's surroundings being the same on each awakening, let them all be different. For example, before each awakening let a unique random number between 1 and 1,000,000 be chosen and stenciled on the wall. The number carries no information about whether the coin came heads or tails or what day it is. And let a photograph, showing the number, be taken of each awakening. On the back of each photograph, a notation is made of whether the coin came heads or tails and what day it was.

Now we can play exactly the same game with *anyone*, just using the photographs. We can ask, for example, what they would offer as fair odds (having been informed of the entire protocol) for the notation on the back of the photograph with the number 578,534 to read "heads", and the relevant photograph can be identified in a perfectly qualitative way. By the very same reasoning that Sleeping Beauty employs, that person will also arrive at 1:2 as fair odds, because (about) 1/3 of the photos have "heads" written on the back and 2/3 have "tails". At no point is any irreducibly *de se* proposition introduced or conditionalized on. The entire problematic is adequately described in entirely *de re* terms, about how to bet on a card with a certain unique visible number on it.

The same sort of analysis, though, can be applied to a purely indexical *de se* proposition. Each time Sleeping Beauty is awoken between Sunday and Wednesday, she can ask herself what day it is (an indexical question). And also, she can be asked to place fair odds on the proposition that it is, at that moment, Monday rather than Tuesday.

As far as the fair odds go, the analysis is the same. If the experiment is repeated 1,000 times and the coin is fair, then exactly 1,000 of the awakenings will be Monday awakenings and approximately 500 will be Tuesday. So the fair odds are 1: ½ (i.e. 2:1). The person betting on Monday has to put in twice as much to the pot as the person betting on Tuesday.

What about her subjective credence that it is Monday? In normal circumstances, the way we come to the conclusion that it is Monday upon awakening is this: we find that the last thing we remember is going to sleep Sunday night and know that we almost always sleep less than a day and remember what happened the last time we woke up. But that reasoning is not available to Sleeping Beauty on account of the possibility of having been given a drug to erase her memory of

waking on Monday. On the other hand, since she knows the experimental protocol she knows the chances of that. So she reasons that there is a 50% chance that the coin came heads and she never got the drug, and a 50% chance it came tails. In the latter case, it is equally likely that it is Monday as Tuesday since there will be an equal number of Monday and Tuesday wakings. So she has a .75 credence it is Monday and .25 it is Tuesday.

Just as in betting on heads or tails, her fair betting odds diverge from her subjective credence. The credence is that of a halfer solution (because she regards the coin as fair) but the betting resembles that of a thirder. And the reason for the divergence, once again, lies just in the peculiarity of the situation. Everyday reasoning won't work because it is not an everyday circumstance. But given the information she has, that is the only rational way to proceed.

And clearly no issues about weird fissioning of persons ever arises, or anything even in the neighborhood.

So like the St. Petersburg puzzle, Sleeping Beauty seems to pose no difficulty at all once the proper method for practical decision-making is adopted. The reason the resolution has been elusive is the usual one: people were looking in the wrong place. The issue is not indexicals, or *de se* belief, or conditionalization: it is how to approach practical reasoning. And because of the habit of trying to derive subjective credences from views about fair odds, incorrect conclusions about Sleeping Beauty's rational credences were drawn. She should always *believe* that the coin was equally likely to come (or have come) heads as tails. But she should always *bet* as she would in a normal situation for a coin biased toward tails. These propositions are not in any conflict or tension because her situation is not normal. The reason the case has seemed puzzling is simply this: people are tempted to arrive at the rational credences by back-formation from the rational betting behavior, using the Principal Principle and the expected value rule. Once one rejects the latter, everything falls into place.

Bibliography

Albert, David *Time and Chance*. Harvard University Press, 2000

Benci, Vieri, Leon Horsten and Sylvia Wenmackers “Infinitesimal Probabilities”. *Brit. J. Phil. Sci.* **69** (2018), 509–552

DiBella, Nicholas “Qualitative Probability and Infinitesimal Probability”, unpublished manuscript, version of September 7, 2018

Fowler, D. H. *The Mathematics of Plato’s Academy: A Reconstruction*. Clarendon Press, 1987

Hájek, Alan “Interpretations of Probability”, *Stanford Encyclopedia of Philosophy*. downloaded Oct. 20, 2020

Hájek, Alan “Staying Regular”, Ms.

Hawthorne, James “A Logic of Comparative Support: Qualitative Conditional Probability Relations Representable by Popper Functions”. In Alan Hájek and Chris Hitchcock, eds., *Oxford Handbook of Probabilities and Philosophy*, Oxford University Press, 2016

Maudlin, Tim *New Foundations for Physical Geometry: The Theory of Linear Structures*. Oxford, 2014

Maudlin, Tim *New Foundations for Physical Geometry: Full Discrete Geometry* (MS in progress)

von Plato, Jan “The Method of Arbitrary Functions”, *BJPS* **34**, March 1983