

Alan Hájek

Philosophy Program

Research School of the Social Sciences

Australian National University

alan.hajek@anu.edu.au

A Puzzle About Degree of Belief

ABSTRACT

Fill in the blank:

Truth is to belief, as ___ is to degree of belief.

Orthodox Bayesianism requires exactly this much of you: your credences should be *coherent*—conform to the probability calculus—and they should change under the impact of new evidence by conditionalizing on that evidence. The scandal of orthodox Bayesianism is its extreme permissiveness, sanctioning as it does credences that are radically out of step with the world. We are hardly so tolerant when it comes to all-or-nothing beliefs: it is generally acknowledged that *truth* is a virtue (perhaps the most important one) that a belief may or may not have. What is the analogous virtue for a degree of belief?

Van Fraassen, Lange and others would fill in the blank by appeal to the notion of *calibration*, a measure of the extent to which degrees of belief track corresponding relative frequencies. I marshal several arguments against this answer. My own answer is:

agreement with objective chance.

Some will complain that the notion of chance is mysterious, or even nonsense. I reply by pointing out several things that we know about chance. The centerpiece of my argument is a ‘magic trick’: Give me any object, any number between 0 and 1 inclusive, and a specified accuracy, and I will use the object to generate an event whose chance is the number given to the accuracy specified.

A Puzzle About Degree of Belief

1. The puzzle

Fill in the blank:

Truth is to belief, as ___ is to degree of belief.

Let us begin by distinguishing *belief* on the one hand, and *degree of belief* (or *credence*) on the other. One tradition in epistemology, associated with Descartes, regards opinion as an all-or-nothing matter: it does not come in intermediate degrees. So the only doxastic categories on this view are outright belief and disbelief (and, if you like, suspension of belief). Bayesian epistemology, however, rightly sees opinion as far more nuanced than that. I don't outright *believe* that it will rain today, but I am considerably more confident that it will rain than not. It seems that my opinion is better represented probabilistically—I assign a credence of 0.8, say, for rain. Likewise for various other judgments of mine (I daresay the majority of them), that fall short of total belief, but that nevertheless have a certain structure, one that is apparently captured by the probability calculus.

Orthodox Bayesianism requires exactly this much of you: your credences should be *coherent*—conform to the probability calculus—and they should change under the impact of new evidence by conditionalizing on that evidence. The scandal of orthodox Bayesianism is its extreme permissiveness.¹ For example, you are beyond its reproach if you assign probability 1 to a turkey proving Goldbach's conjecture—provided you give probability 0 to a turkey *not* proving the conjecture, your other degrees of belief are also coherent, and you arrived at that state of mind by a lifetime of conditionalizing. Slightly less orthodox Bayesians urge you also to stay *regular*: to assign probability 1 only to logical truths and probability 0 only to logical contradictions. This yields other unhappy

¹ Cf. Kyburg (1968).

results: while you may be beyond reproach for assigning 0.9999 to a turkey proving Goldbach's conjecture, you are denounced for assigning it the seemingly more sensible value of 0. Orthodox Bayesians, and their slightly less orthodox brethren who advocate regularity, have no way of recognizing various probability assignments as *defective*, radically out of step with the world. Coherence rules, OK?²

We are hardly so tolerant when it comes to all-or-nothing beliefs. Suppose that Pierre fully believes that L.A. is on the moon, while also fully believing that it is not in the U.S.A.. These beliefs are deductively consistent, to be sure; but consistency is cheap. We clearly hold beliefs to a higher standard than that, for we regard Pierre's as patently defective. What is wrong with them, of course, is that they are *false*: they do not 'correspond to' the facts about L.A.'s location, they do not 'fit' the way things actually are. And it is a desideratum of a belief—perhaps the most important one—that it be *true*. For short, beliefs are governed by a norm of *veracity*.

² A smokescreen that often obscures clear vision of this problem derives from the so-called "convergence theorems" that some Bayesians tout (see Earman 1992 for discussion). These theorems show, roughly, that in the long run, agents who give positive probability to all possibilities, and whose stream of evidence is sufficiently rich, will eventually be driven (by repeated conditionalization on the evidence) arbitrarily close to certainty of the truth regarding the world in which they live, and arbitrarily close to agreeing with each other. One problem is that the theorems say nothing about how quickly these convergences occur. In particular, they do not explain the unanimity that we often reach, and often rather rapidly. Indeed, for any range of evidence, we can find in principle an agent whose prior is so pathological that conditionalizing on that evidence will not get him or her anywhere near the truth, or the rest of us. Furthermore, the theorems do not give iron-clad guarantees that there will be the desired convergences, but only promise them with the disclaimer '...with probability one'. The concern is that the probability here is that of *the agent himself or herself*. Note, moreover, that Bayesians who, like van Fraassen (1989), eschew updating rules as rational requirements cannot appeal to such convergence results.

In any case, the condition that the priors not zero out genuine possibilities (for no amount of conditionalizing can raise a zero probability assignment) is not as innocuous as Bayesians might like to think. Since there are in principle uncountably many mutually exclusive hypotheses, any probability measure zeroes out uncountably many of them, thus stacking the deck infinitely in favor of some hypotheses over others. A gesture to infinitesimals at this point will not suffice. For the convergence theorems are premised on Kolmogorov's axioms holding. It needs to be proven that analogous theorems hold when we move to probability functions that take non-standard values. In particular, Kolmogorov's axiom of countable additivity cannot even be *stated* in the language of non-standard analysis. See Hajek (2003b) for further reservations about the use of infinitesimals.

And even if we all come together in the Peircean limit, the Bayesian must allow that we disagree with each other on many things *now*, and will at all times. There is surely some fact of the matter of who among us, if any, is right at a given time. All coherent agents are equal—but some are more equal than others.

Coherence is likewise cheap. Just as our beliefs aim for more than consistency, our credences surely aim for more than coherence. What does a credence's 'corresponding to' or 'fitting' the world amount to? What plays the role for subjective probability analogous to the role that truth plays for all-or-nothing belief? Presumably *something* plays this role, for probabilistic judgments seem also to be governed by a norm of veracity. How else can we explain the fact that we regard some weather forecasters as better than others, even when they couch their predictions probabilistically? When different people with the same evidential base assign different probabilities to the same event, we think that they are *disagreeing* about something. Faced with such disagreement, the parties involved can reasonably argue for the superiority of their own probability assignments. This would make no sense if coherence were the yardstick by which such opinions were judged. I claim that there is another yardstick—the probabilistic analogue of truth—and according to it they cannot both measure up.

Of course, beliefs may have various other virtues beyond being true. They may be justified, informative, useful, comforting, and what have you. Degrees of belief may have similar virtues. Those are not the topic of this paper; I am focusing exclusively on the analogue of truth for degrees of belief. And just as a belief may be unjustified, uninformative, useless, discomfoting, and yet still be true, so a degree of belief may lack various virtues, yet still have the analogue of truth. In particular, just as a given belief is *vindicated* by being true, however it fares by the other criteria, so too a given degree of belief may be vindicated, however it fares by the other criteria. A belief is vindicated by truth; a degree of belief is vindicated by ____?³

³ I am assuming, of course, that there *are* such things as beliefs, truth, and degrees of belief. Churchland (1981) is skeptical of the first of these assumptions; a Nietzschean is skeptical of the second; Harman (1986) is skeptical of the third.

2. Perfect calibration and potential perfect calibration

Van Fraassen (1983, 1984) and Lange (1999) apparently would fill in the blank by appeal to the notion of *calibration*. Indeed, in a discussion of evaluating the probabilistic announcements of weather forecasters, van Fraassen (1983) distinguishes probabilistic informativeness (roughly, closeness of the probabilities to 0 or 1) and calibration, regarding them respectively as the counterparts of informativeness and truth in non-statistical theories. He concludes: “Calibration plays the conceptual role that truth, or empirical adequacy, plays in other contexts of discussion” (301). So what is calibration? Here’s how van Fraassen (1984) explains it:

consider a weather forecaster who says in the morning that the probability of rain equals 0.8. That day it either rains or does not. How good a forecaster is he? Clearly to evaluate him we must look at his performance over a longer period of time. Calibration is a measure of agreement between judgments and actual frequencies... This forecaster was perfectly calibrated over the past year, for example, if, for every number r , the proportion of rainy days among those days on which he announced probability r for rain, equalled r . (245)⁴

For instance, to be perfectly calibrated, the proportion of rainy days, among days on which he announced a probability of 0.8 in the morning, must be 0.8; and likewise for all other probabilities that he announces over the course of the year. Van Fraassen goes on to suggest that perfect calibration is a desideratum of probabilistic judgment:

Although perfect calibration may not be a reasonable aim by itself, and hardly to be expected at the best of times, it certainly looks like a virtue. It would seem to

⁴ Strictly speaking, this characterization of the forecaster’s being perfectly calibrated is too weak. The sufficient condition, given in the final sentence, is met by a wildly incoherent forecaster whose probabilities for *rain* match the corresponding relative frequencies of rain, as van Fraassen wants, but who also assigns each morning numbers to *not-rain* completely at random. In that case we should say that his probabilities are perfectly calibrated *with respect to rain*, but not *with respect to not-rain*. The relativization to an event-type is crucial. If, instead, we simply looked at all occasions on which the forecaster gave a probability assignment of r (to some event or other), and we checked whether the relevant events occurred on a proportion r of those occasions, perfect calibration would be all too easy. All the forecaster would have to do is give probability 1/2 for rain and probability 1/2 for not-rain every day of the year. Of the 365 x 2 events to which he assigns probability 1/2, exactly 365 of them occur, since on each day exactly one of ‘rain’ and ‘not-rain’ occurs! (I thank David Dowe for alerting me to this point. It is also made by Lad 1996.) More generally, if there are n possible outcomes for each trial, and m trials, one could trivially guarantee perfect calibration by assigning $1/n$ to each possible outcome: of the mn possible outcomes that receive that value, m of them occur. This new argument for the principle of indifference is surely too good to be true.

be irrational to organize your degrees of belief in such a way as to ruin, a priori, the possibility of perfect calibration.

This surely has *prima facie* plausibility. I take Shimony (1988) and even Ramsey (1926) to be further weighty authorities in its favor. Compare a familiar point about beliefs: it would be irrational to organize your beliefs so that they are inconsistent, since that ruins, a priori, the possibility that they are all true. This shows that we do regard truth of our beliefs to be a desideratum; perfect calibration, then, is offered as the analogous desideratum of our degrees of belief.⁵ Van Fraassen (1983, 1989) also uses the language of ‘vindication’ of degrees of belief and writes: “Vindication I shall explicate in terms of *calibration*.” (1983, 300).

The quote from van Fraassen (1984) continues:

A few qualifications must at once be introduced: this forecaster would not have been perfectly calibrated over the past year if he had announced irrational numbers, or even numbers not equal to some fraction of form $x/365$.

This is not quite right. The forecaster *could* have been perfectly calibrated over the past year even if he had announced numbers not equal to some fraction of the form $x/365$ —indeed, even if he announced *only* such numbers. To see this, suppose that he announced a probability of $1/2$ on the first two days of the year, and $1/363$ every day thereafter; and suppose that it rained on only the first and last days of the year. Then he was perfectly calibrated, even though neither $1/2$ nor $1/363$ is of the form $x/365$. But van Fraassen is certainly right that someone who assigns irrational probabilities cannot be perfectly calibrated (assuming that the number of trials is finite, as it is in the weather example).

⁵ In a similar vein, Lange (1999) writes:

In assigning 80% confidence to the claim that it will rain tomorrow, the forecaster is perfectly calibrated if and only if it rains on exactly 80% of the days that are relevantly like tomorrow. If it rains on more than 80% of those days, then the forecaster’s subjective degree of belief is too low; if it rains on less than 80% of those days, then her opinion is too high. If she is perfectly calibrated, then her opinion is right (311).

Also: “if we have shown ourselves entitled to believe a probability distribution to be perfectly calibrated, then ... we have shown ourselves entitled to believe those degrees of confidence to be *right*...” (317, italics in original).

He then finesses the notion of calibration, and correspondingly the aim of a forecaster's subjective probability assignments:

So the only possibility that we should require him not to ruin beforehand is that of arbitrarily close approximation to perfect calibration if he were asked sufficiently often about events that he considers exactly similar to those he was actually asked about during the evaluation period.

Let's follow van Fraassen in calling the weather forecaster *potentially perfectly calibrated* if his hypothetical probabilistic judgments could agree, to arbitrarily close approximation, with corresponding relative frequencies in a suitable hypothetical sequence of trials. (See also Dawid 1982 for a similar notion of calibration that has been influential in the statistics literature.) We then come to one of van Fraassen's main conclusions:

It can now be proved that satisfaction of this criterion of potential perfect calibration is exactly equivalent to satisfaction of the probability calculus...

This gives us another putative source for interest in calibration: it gives us, as van Fraassen says, "a sort of dual to the Dutch Book argument" (245). Much as the consistency of a set of beliefs is equivalent to the possibility of their joint truth, the coherence of a set of degrees of belief is equivalent to their being potentially perfectly calibrated. We will return to this point shortly.

3. Against perfect calibration and potential perfect calibration

Calibration is not always as virtuous as it appears; still less is it *the* virtue that we are seeking, the one that corresponds to *truth*. The desire to improve calibration can be an incentive to ignore, or even to go against, all of one's evidence. Consider a five-day evaluation of a forecaster who will be handsomely rewarded if he is perfectly calibrated. Suppose that his assignments for rain on the first four days are $2/3$, $2/3$, $1/2$, and $1/2$, and that it rains on the first three days and not on the fourth. Then whatever his real opinion about the fifth day, he should announce probability $2/3$ for rain and hope for not-rain! He

can ascertain this from his armchair on the fourth night, without paying an iota of attention to indicators of the weather. Clearly, truth is nothing like perfect calibration in this respect: if a forecaster is rewarded for having a *true belief* about the weather, he should use his total evidence.

When a calibration-driven forecaster announces a probability of x/n (where x and n have no common factors), he is *a priori* committed to announcing that probability at least another $n-1$ times, and in any case some multiple of n times in total. For example, if he announces probability $3/10$ for rain on one day, then he thereby forces his own hand to announce that probability at least 9 more times, and in any case some multiple of 10 times in total. All this, again, irrespective of indicators of the weather, or his true opinions. It also seems odd that the more *unlikely* you find rain on a given day, the more times you should feel committed to *repeating* that assignment on subsequent days; but that is the road along which the goal of perfect calibration leads you. In assigning probability $1/10$ on a given day, you are committed to repeating that assignment at least 9 times thereafter; in assigning probability $1/20$ on a given day, you are committed to repeating that assignment at least 19 times thereafter; and so on. And having done so, the cycle starts again. Similarly, the goal of perfect calibration often sets up strange incentives to ‘round off’ your probability assignments, and thus to dissimulate, in certain ways. Suppose that your probability for rain tomorrow is in fact 0.417. To be perfectly calibrated, you would need to assign that probability some multiple of 1000 times, something that you cannot do in an evaluation period of 365 days, for example. On the other hand, an assignment of 0.4 need only be given 5 times, so you would in this sense do better rounding off your genuine assignment. Still, giving even 5 such assignments may well be an unwelcome commitment to undertake (especially if you do not think that the weather distribution is stationary over time). Rounding off again to $1/2$ eases the burden to just a single forced repetition; and rounding off yet again to 0 or 1 relieves you of the burden altogether! These are all consequences of the fact that perfect calibration is

not a property of a *single* probability assignment, but rather a global property of a *set* of such assignments.

Familiar arguments against frequentism—the thesis that chances are relative frequencies within a suitably chosen reference class—return to haunt calibration. Various events with apparently intermediate probabilities belong naturally only to single-membered reference classes—the ‘problem of the single case’. A cosmologist who assigns intermediate probability to the universe being open ruins, a priori, the possibility of perfect calibration, but he or she need not be faulted for that.⁶ The same is true of anyone who assigns irrational-valued probabilities to anything (for plausibly no actual reference class is infinite). And when we ask how good is our weather forecaster’s probability assignment of 0.8 to rain *today*, our question surely does not concern what he says and what happens on *other* days. The refinement to potential perfect calibration leads to further problems, again familiar from the demise of a counterpart version of frequentism, one that identifies chances with limiting relative frequencies in a hypothetical infinite sequence of trials (see, e.g., von Mises 1957). There is simply no fact of the matter as to what a hypothetical infinite sequence of trials of the relevant sort would be⁷, hence no fact of the matter of how well our forecaster’s (equally hypothetical) assignments track the corresponding limiting relative frequencies. The limiting relative frequency of a given attribute can be changed to any value one wants by suitably reordering the trials. And one wonders what relevance an *imaginary* sequence of outcomes has to our assessment of the forecaster as he *actually* is, especially when it seems that any world that instantiates an *infinite* sequence of trials of the requisite kind would have to be very different from the actual world, presumably differing even on the laws of nature. (Think what it would take for the coin in my pocket, which has only been

⁶ In fact, proponents of regularity *require* it of him or her!

⁷ Indeed, there is not even a fact of the matter as to what a hypothetical *finite* sequence of trials of the relevant sort would be, nor even as to what a *single* hypothetical trial would be. Today turned out to be sunny; but how would the first hypothetical day have turned out? Cf. Jeffrey (1977).

tossed once, to be tossed infinitely many times—never wearing out, and never running short of people willing to toss it!)

Potential perfect calibration inherits these frequentist problems, and then it adds more of its own. In van Fraassen’s example, it depends on which days the weather forecaster *himself* “considers exactly similar” to which, suggesting that whether or not he is perfectly calibrated is in the eye of the beholder—in this case, himself. Van Fraassen elsewhere (1989) is skeptical of *anti-nominalism*, the thesis that nature comes pre-carved with objective respects of similarity. The worry is that any probability assignment whatsoever can be given the ‘vindication’ of potential perfect calibration by considering suitably ‘unnatural’ or ‘grue’-like respects of similarity. But even this assumes that we can make sense of the notion of “*exact* similarity”, yet it strikes me as an oxymoron—it is neither numerical identity, nor mere similarity (which is necessarily *inexact*). I think that the turn of phrase is revealing, symptomatic as it is of a tightrope that potential calibration has to walk. Exact copies of a given day will not do, for then all relative frequencies are trivialized. But inexact copies of a given day are problematic too, since they may differ from it and among themselves in respects that affect the probability of rain.

There are in general many conflicting ways to be perfectly calibrated to a given finite set of outcomes. Indeed, every partition of the set—even every ‘unnatural’ partition, if there are such—yields a way to be perfectly calibrated: across each cell of a given partition, assign the relative frequency of the outcome of interest in that cell. In our earlier example, suppose it did not rain on the fifth day, so that in assigning probabilities $\langle \frac{2}{3}, \frac{2}{3}, \frac{1}{2}, \frac{1}{2}, \frac{2}{3} \rangle$, our forecaster was perfectly calibrated to the actual outcomes, $\langle \text{Rain}, \text{Rain}, \text{Rain}, \neg \text{Rain}, \neg \text{Rain} \rangle$. Let D be the set of days, which we will represent as $\{1, 2, 3, 4, 5\}$. One partition of D is $\{\{1,2,5\}, \{3,4\}\}$; it corresponds to our forecaster’s assignment. Another partition of D is the set of D ’s singletons; it corresponds to the maximally informative, perfectly calibrated assignment $\langle 1,1,1,0,0 \rangle$. Another partition of

D is the singleton of D itself; it corresponds to the maximally uninformative, perfectly calibrated assignment $\langle \frac{3}{5}, \frac{3}{5}, \frac{3}{5}, \frac{3}{5}, \frac{3}{5} \rangle$. And so on. We have many perfectly calibrated, yet disagreeing, assignments. Again, there can be no analogue of this for truth: no two disagreeing beliefs can both be true.

There are still more conflicting ways of being potentially perfectly calibrated. Recall van Fraassen's conclusion: potential perfect calibration is equivalent to satisfaction of the probability calculus; that is, it is equivalent to coherence. But as I have said, coherence is cheap. There are infinitely many ways of coherently assigning degrees of belief to 'rain' and 'not-rain'; each one, then, corresponds to a way of being potentially perfectly calibrated. Despite their disagreements, all coherent forecasters are equally good by the lights of potential perfect calibration. So we do not get a non-trivial answer to the question of how good was van Fraassen's forecaster who assigned 0.8 to rain that day: as long as he was coherent, he was no better and no worse than any other coherent forecaster.

Potential perfect calibration may well fill in the blank in a *different* puzzle:

Consistency is to a set of beliefs, as ___ is to a set of degrees of belief.

But we still have not found a satisfactory answer to our original puzzle.

4. For chances

Here, then, is my answer:

Truth is to belief, as agreement with objective chance is to degree of belief.

The forecaster's degree of belief of 0.8 for rain that day was vindicated iff the objective chance of rain that day, at the time at which he made the announcement, was 0.8. More generally: A degree of belief for a proposition that agrees with the corresponding objective chance for that proposition, at the time at which the degree of belief is held, has the virtue that that we have been seeking: the probabilistic analogue of truth. And the

closer a degree of belief is to the corresponding objective chance, the closer it is to having that virtue.

This requires us to recognize chance as an objective, mind-independent feature of the world, much as we take mass, charge, and so on as objective, mind-independent features of the world. And much as one's subjective estimate of the mass of an object 'aligns with the truth' to the extent that it approximates the objective mass of that object, so one's subjective probability of an event 'aligns with the truth' to the extent that it approximates the objective probability of that event.

The arguments against calibration will not trouble us. On the fifth day of his evaluation period, our forecaster has no incentive to ignore his evidence if he wants to align with the relevant chances. How good his assignment is that day does not depend on what he says or what happens on other days. He is not under pressure *a priori* to repeat a given probability assignment any number of times, nor has he any incentive to dissimulate. Reordering the trials makes no difference to the chances of propositions concerning them. Only chances in the actual world are relevant. There is a fact of the matter about what the chances of propositions are, so there is a fact of the matter of how well people's assignments to these propositions are matching them. (Of course, the chances may be difficult to ascertain—though, see section 6, below—but that's beside the point.) Plausibly, agreement with a corresponding chance is a property that a single probability assignment can have, rather than being a global property of a set of such assignments (so the assignment of 0.8 to rain has the probabilistic analogue of being true, or not, in isolation). And there is only one way to match a given objective chance; no two disagreeing assignments can both agree with it. Mere coherence is not enough.

It isn't just that in all these ways, chance-tracking is superior to perfect calibration as a desideratum for credences (although that too is true). More than that: in all these ways, chance-tracking is a desideratum for credences *analogous to truth-tracking for beliefs*. A forecaster who each morning assigns a *truth-value* to it raining or not raining that day

should use his total evidence. The truth-value of what he says about a given day does not depend on what he says or on what happens on other days. He is not under pressure *a priori* to repeat a given truth-assignment any number of times, nor has he any incentive to dissimulate. Reordering the trials makes no difference to the truth-values of propositions concerning them. Only truth-values in the actual world are relevant. There is a fact of the matter of what the truth-values of propositions are, so there is a fact of the matter of how well people's assignments to these propositions are matching them. (Of course, the truth-values may be difficult to ascertain, but that's beside the point.) Plausibly, truth is a property that a single belief can have, rather than a global property of a set of beliefs (so the belief that it will rain today is true, or not, in isolation). And there is only one way to match a given truth-value; no two disagreeing assignments can both agree with it. Mere consistency is not enough.

We may bring out the analogy between truth-value-matching and chance-matching further by considering Lewis' (1980) famous *Principal Principle*. Let ' P ' be the subjective probability function of a rational agent, let A be a proposition, let ' ch_t ' be the chance function at a time t , and let E be any proposition that is admissible at t (roughly: that does not yield information taken to be relevant to the truth of A beyond what its chance is). The Principle is:

$$P(A \mid (ch_t(A) = x) \ \& \ E) = x \text{ for all } x \text{ such that } P((ch_t(A) = x) \ \& \ E) > 0.$$

For example, my degree of belief now that this coin toss will land heads, given that its chance now of landing heads is $3/4$, is $3/4$. It follows that if the agent is certain of the chance of A , her credence will follow suit. A similar principle governs all-or-nothing beliefs. Let ' B ' be the *belief function* of our rational agent: $B(X) = 1$ if the agent believes X , and $B(X) = 0$ otherwise. Let $B(X|Y) = 1$ if the agent believes X given the information that Y , and $B(X|Y) = 0$ otherwise. And let T be the *truth function*: $T(X) = 1$ if X is true, and $T(X) = 0$ otherwise. Then we may introduce the *Truth-Tracking Principle*:

$$B(X \mid T(X) = 1) = 1.$$

Given the information that X is true, the rational agent will believe X . Failure to do so is not inconsistent, but it does give rise to Moore-paradoxical sentences of the form ‘ X is true, but I do not believe that X ’. Similarly, failure to adhere to the Principal Principle is not probabilistically incoherent (it does not violate the probability calculus or expose the agent to a Dutch Book), but it does give rise to Moore-paradoxical-ish sentences of the form ‘ A has chance x , but my credence in A is not equal to x ’. My account of chance-matching as the analogue of truth-value-matching explains these phenomena.

What it doesn’t explain, I admit, is the role that admissibility plays in the Principal Principle, for there is no analogue of that in the Truth-Tracking Principle. After all, *inadmissible* information would undermine the link between chances and credences. For example, $P(\text{Heads} \mid (ch_{now}(\text{Heads}) = 1/2) \ \& \ \text{Heads}) = 1$, not $1/2$. To be sure, inadmissible evidence is very hard to come by (and the cases discussed typically involve far-fetched stories of time-travelers, crystal balls, and the like). But there is simply no such thing as inadmissible information for our Truth-Tracking Principle: $B(X \mid (T(X) = 1) \ \& \ Z) = 1$ for all Z (compatible with $T(X) = 1$). Conditional probability assignments are *non-monotonic* in the sense that they can be decreased by conjoining additional information to the conditions; but conditional *belief* assignments are in this sense monotonic. Does this undercut my claim that chance is the probabilistic counterpart of truth? I think not. We already knew that while deductive logic is monotonic (in the related sense that a valid argument is never rendered invalid by the addition of a premise), inductive logic is non-monotonic. Put simply, when it comes to probabilities, non-monotonicity is part of the territory. Whatever the probabilistic counterpart to truth is, we should expect it to be non-monotonic. I am thus not troubled that my proposal for the counterpart is; in fact, I am pleased that it is.

My viewpoint stands in contrast to more ‘pragmatic’ approaches to the concept of chance, as advocated for example by Skyrms (1984) and van Fraassen (1989), according to which chances correspond to certain features of subjective probability functions, but

are not out there in the world. Now, I am certainly not the first to pledge allegiance to a robust notion of objective chance—see, for example, Popper (1959), Mellor (1971) and Lewis (1980). I *am* offering a new reason for such allegiance. We should believe in chances because of the norm of veracity for degrees of belief. *Chances are the vindicators of credences.*

They are likewise the vindicators of probabilistic statements made by *probabilistic theories*, and thus of the theories themselves. Some of us still believe that science aims to give us true theories about the world. Even if it doesn't, surely scientific theories can *be* true, as even many a scientific anti-realist will admit. What does this amount to in the case of a probabilistic theory? I submit: the probabilities that it assigns agree with the corresponding chances.

Moreover, conditional chances are the vindicators of conditional credences, and of statements of conditional probability made by probabilistic theories. My credence in rain tomorrow, given that it rained yesterday, is vindicated just in case it matches the corresponding conditional chance. Quantum mechanics tells us that the probability that a certain electron entering a Stern-Gerlach apparatus will be measured to be spin 'up', given that it is measured for spin (in a particular direction), is $1/2$. This claim is vindicated just in case the corresponding conditional chance is $1/2$. And the entire theory of quantum mechanics is vindicated just in case all such probabilistic claims that it makes agree with the corresponding objective chances.

Some authors emphasize that chances must be relativized to an experimental arrangement (Popper 1959), or to a chance set-up, (Giere 1973), or to a kind of trial (Levi 1990). While there may be some finessing of the details, the common idea here is that chance is really a function of two propositions: let's call them an *outcome* proposition, and a *context* proposition. This is tantamount to regarding all chances as conditional (although sometimes context will make one condition so salient that we may assume it without explicit mention; see Hájek 2003a). So we might have: the chance that this coin

lands heads *given* that it is tossed high and with a vigorous spin is $1/2$, while the chance that it lands heads *given* that it is initially facing heads up, and feebly released a centimeter above the floor, is 1... In any case, these conditional chances will be the vindicators of corresponding credences with the same outcome propositions and context propositions.

I can remain somewhat neutral on just how much out there is chancy. At one extreme, we have the view that every proposition has a chance value—I take this to be Lewis’ view. At the other extreme, we have the view that no proposition has a chance value—this is the view of de Finetti and van Fraassen, as we will shortly see. In between these two extremes there is a spectrum of positions. Cartwright (1999), for example, thinks that chances exist, but only for certain very special kinds of set-ups (what she calls “nomological machines”). For what it’s worth, I would locate my own view down the profligate, chance-rich end of the spectrum (while granting that there may well be some chance-gaps). But my position in this paper is tenable as long as at least *some* propositions have chances, and so it is compatible with more austere views. When there are chances out there, they are the vindicators of corresponding credences. Compare: I can remain somewhat neutral on just how much out there has a *truth-value*. We have a spectrum of positions, from the view at one extreme that every proposition has a truth-value, to the view at the other extreme that no proposition has a truth-value. For what it’s worth, I would locate my own view down the profligate, truth-value-rich end of the spectrum (while granting that there may well be some truth-gaps). But my position in this paper is tenable as long as at least *some* propositions have truth-values, and so it is compatible with more austere views. When there are truths out there, they are the vindicators of corresponding beliefs.

I have given my answer to the puzzle—I have filled in the blank. But as always, there are objections to face.

5. Objections

“The world might have extreme chances”

Suppose that we live in a world in which all chances are 0 or 1. A Laplacean thinks that this is not merely a supposition; in any case, even in this post-quantum mechanics era it is still a live possibility. Yet perfectly rational credences, and the probability assignments of our best theories, will still often be intermediate. The first objection is that this shows that my answer to the puzzle cannot be right.

I reply: In such a world, the only probabilities that are vindicated *are* 0's and 1's, be they assigned by an agent or by a theory. Of course, it might not be rational for you to make such assignments, or to accept such a theory, given the evidence at your disposal—doing so may not be *justified*. (The analogous point regarding all-or-nothing beliefs is a commonplace.) But I am not interested here in *that* desideratum of probabilistic judgments. To repeat, I am confining myself to finding the probabilistic analogue of truth. And I claim to have found it in conformity to chance.

“The world has enclaves with extreme chances”

But let us pursue the line of objection further. Even if not all chances in the world are extreme, still there are surely enclaves in which they are. Perhaps coin tossing, for example, is ‘deterministic’ (in this sense), even if the underlying microphysics is not. There are doubtless times when we find ourselves in such an enclave without realizing it, or without knowing which extreme is realized where. Our situation is often rather like this: the coin that I am about to toss is either two-headed or two-tailed, but you do not know which. What is the probability that it lands heads? Very reasonably, you assign a probability of 1/2, even though you know that the chance of heads is either 1 or 0. So it is rational here to assign a credence that you *know* does not match the corresponding

chance. Truth and belief, the objection concludes, do *not* behave like this: it is never rational for you to adopt a belief that you *know* is not true.⁸

I reply: It *is* sometimes rational for you to adopt a doxastic state that you know is not true—namely, a state of *suspension of belief*. What all-or-nothing belief should you adopt regarding the outcome of the coin toss? None—or if you like, just the tautologous belief ‘the coin will either land heads or not’. You are not rationally entitled to any stronger belief regarding the outcome. In particular, you should suspend belief in the coin landing heads and suspend belief in the coin landing tails. But in doing so, *you preclude the possibility of having a true belief regarding the outcome*. And praise be to you, I say. That just shows that truth is not the only epistemic desideratum—something I said at the outset.

Suppose that the coin is in fact two-headed. Someone who shares your impoverished evidential state, but who nonetheless impulsively believes that the coin will land heads, is getting something *right*—is in step with the world in one respect—even if he or she deserves no credit for that. (Here ‘credit’ has justificatory overtones.) When you suspend belief in the outcome, you are *failing* to get something right—failing to be in step with the world in one respect—even if you deserve no blame for that. (Again, ‘blame’ with justificatory overtones.) Similarly, someone who shares your impoverished evidential state, but who impulsively assigns heads probability 1 nonetheless, is getting something *right*—is in step with the world in one respect—even if he or she deserves no credit for that. When you assign heads probability 1/2, you are getting something *wrong*—are out of step with the world in one respect—even if you deserve no blame for that.

I am prepared to concede that the analogy between all-or-nothing belief and credence is not perfect; but one should not expect it to be perfect. I have already noted that one side of the analogy involves ‘non-monotonic’ quantities, while the other involves

⁸ Note that this would equally be an objection to perfect calibration being the analogue of truth—for however many times we toss the coin, you *know* that your credence will not match the relative frequency of heads.

‘monotonic’ quantities. Moreover, one side of the analogy involves *bivalent* quantities, while the other side involves *continuous* quantities. Whatever it is that fills in the blank—the vindicator of credences—it must come in uncountably infinitely many degrees, for credences do. Again, I submit: chances are right for the job.

“Nothing fills in the blank”

You may say that there is *no* probabilistic analogue of truth: nothing fills in the blank in the puzzle. Intermediate degrees of belief, you say, are not made ‘true’ by anything; nothing could vindicate them; there is no norm of veracity for them.

I reply: I find this view quite implausible. If credences ‘float free’ of the world in this way, then why should we ever update them? Why do we think that some updatings are better than others? Why, for example, should we consult our weather forecaster’s judgments? Why should we use our total evidence rather than merely partial evidence, or none at all? Come to think of it, what use are credences at all on this view? It sees a radical disjuncture in our epistemology. After all, we treat all-or-nothing beliefs with rather more respect.⁹ Even false beliefs *could be* vindicated—namely, in a world where they turn out to be true—even though as a matter of fact, they are not. Indeed, I’m inclined to say that even suspensions of belief could be vindicated—namely, by corresponding truth *gaps*, if there are such things. So surely intermediate degrees of belief, which are ‘truer’ than false beliefs, and more committal than suspensions of belief, can be vindicated—and setting aside a certain kind of widespread skepticism, some of our degrees of belief *are* vindicated.

“‘Truth’ fills in the blank”

⁹ Most of us, anyway.

You may say that there is a probabilistic analogue of truth—namely, truth. You say that ideally an agent assigns probability 1 to truths, and probability 0 to falsehoods; credences don't get truer than that.

I reply: That may well be right, but it gets my puzzle back to front. You are starting with the outcomes, so to speak, and stating that credences that align with them are the truest ones. Start instead with a credence in some proposition—say, 0.8—and ask what vindicates *it*, the way the truth of a proposition vindicates a belief in that proposition. Here is a belief in some proposition; it corresponds to the way the world is if the proposition is true. Here is a credence in some proposition; it corresponds to the way the world is if it agrees with the chance of that proposition.

“Chances exist, but agreement with chance does not fill in the blank”

The next objection is related to the last one. You may agree with me that chances exist, but disagree with the way that I call upon them to solve my puzzle. You challenge me with this sort of case:¹⁰ Consider an event that is genuinely chancy—our example will involve radium decay. Now consider the following conversation:

Verity: This radium atom will decay within 1500 years. My credence is 1: I fully believe it.

Fortuna: My credence is 0.5 that this radium atom will decay within 1500 years. Five minutes later, while Verity and Fortuna are still arguing, the atom decays. The conversation continues:

Verity: I was right! My belief of five minutes ago was true!

Fortuna: Sorry—you were wrong and I was right. The analogue of truth for a credence is its agreement with the corresponding objective chance. The objective chance that the particle would decay within 1500 years was 0.5, not 1.

¹⁰ Paul Bartha and Denis Robinson are the real-life objectors who played the role of ‘you’ here. I thank Paul for most of the script in the conversation that follows.

The analogue of falsity for a credence is disagreement with the objective chance, and your credence disagreed with the chance.

The worry here is that by my lights credence has to serve two masters. For full belief is a limiting case of credence: namely, 1.¹¹ *Qua* belief, it is vindicated if the corresponding proposition is true; but *qua* credence, it is vindicated if it agrees with the corresponding chance. When that chance is less than 1, and yet the proposition is true, we get conflicting verdicts. By my lights, the objection continues, any full belief about a chancy proposition is “false” (or the analogue of false); yet the proposition may be in fact true. In short, the norm of veracity for full beliefs conflicts with the norm of veracity for credences.

I reply: This objection assumes that beliefs about chancy matters are determinately true or false. This assumption is controversial—it is controverted by Belnap and Green (1994), Belnap, Perloff and Xu (2001), and MacFarlane (2003) among others. Plausibly, at the time at which Verity first spoke there was an objectively possible future in which the atom decays within 1500 years, and another one in which it does not, so the statement ‘the atom will decay within 1500 years’ was *indeterminate* at that time. Then Verity is not entitled to her boast five minutes later. Fortuna, on the other hand, is entitled to hers, although she is excessively harsh on Verity: the latter was neither ‘wrong’ nor ‘right’.

However, I do not have to take a stand either way on the controversial issue of the determinacy or otherwise of truth-values of beliefs about chancy matters. Indeed, let us grant for the sake of the argument that they are determinate, just as the objection would have it. Whether we like it or not, chancy processes have a dual aspect: their actual outcomes make true certain propositions (and make false others), yet these outcomes are probabilistic. Full beliefs about these processes that are true align with the actual outcomes; intermediate credences that are ‘true’ (or the analogue of true) align with the

¹¹ I understand ‘all-or-nothing’ beliefs to be more inclusive than ‘full’ beliefs. I can be sufficiently confident of something to believe it, without giving it probability 1; in fact, I would say that most of what I believe I assign probability less than one. Indeed, a fan of regularity would insist that rational agents have no *full beliefs* about contingent matters whatsoever, but might still allow that they have *beliefs* about such matters.

chances. Far from being in conflict, these two aspects co-exist happily. In the dialogue, both Verity and Fortuna got something right. Making the case for Verity is easy (once we have granted that her belief was determinately true, as we have), and she did so herself. But Fortuna also deserves credit—indeed, 100% credit. After all, her credence is in perfect agreement with a fact about the world: the half-life of this radium atom. She should just be more ecumenical—she doesn't have a monopoly on the truth, the way she would have us believe. Verity's and Fortuna's credences happily co-exist, because each matches a genuine aspect of the world.¹²

Similarly, a probabilistic theory such as quantum mechanics, and a theory that makes only non-probabilistic claims about the same phenomena (by predicting the actual outcomes), could both be true; indeed, a theory that *combined* all their claims could still be true. The probabilistic theory is true if all its probability assignments match the corresponding chances; the non-probabilistic theory is true if all of its statements are simply true.

There is just one master to be served here: correspondence to the world. But there are both outcomes *and* probabilities in the world to which our doxastic mental states (beliefs, credences), and our theories (non-probabilistic, probabilistic), may or may not correspond. In thinking about my puzzle, I want you to begin with a credence—say, 0.8—and ask what has to be the case for it to correspond to the world. That is the probabilistic analogue of truth.

“True credences are just true beliefs about chance”

¹² The dual aspect of chancy processes is familiar in the seemingly conflicting retrospective evaluations that we sometimes make of decisions, which can also be viewed from the perspective of actual outcomes or of probabilities. Suppose that I decide not to buy fire insurance for my house, knowing full well that there is a serious risk of fire where I live. Soon afterwards, a fire ravages my neighborhood, but by luck it leaves my house untouched. Did I make the right decision? Well, yes and no. As far as the outcome is concerned, yes—I saved my money and, as it turned out, I did not lose my house. As far as probabilities are concerned, no—the chance of fire destroying my house was too high to vindicate my complacency. In fact, there is no conflict here: the two aspects co-exist happily, and each evaluation can be regarded as correct in one sense.

You may say that a credence is just a special kind of (all-or-nothing) belief, namely a belief with probabilistic content. You say that Fortuna's credence of 1/2 concerning the radium decay is simply a belief that the chance of decay is 1/2. And in general, a credence of x in X is simply a belief that the chance of X is x .¹³ In a slogan, degrees of belief are just beliefs-about-degrees (of chance). Thus, there is no *analogy* in the puzzle at all. The 'target' of the putative analogy, degree of belief, is just a species of the 'base', belief. Truth plays the same role throughout. Or so you may say.

I reply: Only some credences can be identified with beliefs about chances. For all we know, Fortuna does *not* believe that the chance of the radium decay is 1/2. She may instead give some credence to various incompatible hypotheses about that chance, and then use the law of total probability to determine finally her credence in the decay. We saw an extreme case of such 'mixing' of probabilities in the example of the coin that was two-headed or two-tailed, but you did not know which. Your credence in 'Heads' was 1/2, that is,

$$\begin{aligned} & P(\text{Heads} \mid \text{two-headed}) \cdot P(\text{two-headed}) + P(\text{Heads} \mid \text{two-tailed}) \cdot P(\text{two-tailed}) \\ &= 1 \cdot \frac{1}{2} + 0 \cdot \frac{1}{2}. \end{aligned}$$

But this was *not* to be identified with a belief that the chance was 1/2; on the contrary, you fully believed that the chance was *not* 1/2!

Or perhaps Fortuna lacks the concept of chance altogether. She may have no beliefs about chances, still less the belief that the chance of decay is 1/2. Nevertheless, she may be uncertain about the world, and this may be reflected in her credences, such as her credence of 1/2 in the decay event. Or perhaps she believes that there are chance gaps, and that the decay event is one of them; still, she assigns it a credence. Or perhaps she believes that chances can be vague, and that the decay event has a vague chance; still, she assigns it a sharp credence. In all of these ways, credences can come apart from beliefs about chances.

¹³ Cf. Harman (1986).

The thesis that they should be identified would also render the Principal Principle *trivially* true. That principle allows us to link credences to beliefs about chances, but the credences and the beliefs are still distinct things.¹⁴ Nobody, not even the Principle's greatest fan, should think that it admits of a one-line proof!

“There is more than one way to fill in the blank”

You may say that there is *more than one* probabilistic analogue of truth. You say that truth bears more than one relation to belief, so there is more than one way to fill in the blank. (I leave it to you to say more about these multiple relations.)

I reply: Even if that is so, it is not a problem for me. Suppose, for the sake of the argument, that truth bears numerous relations R_1, R_2, \dots, R_n to belief, and it is only the R_k relation that chances bear to corresponding credences. Still, even if we say no more about the R_k relation, we have the result that chances exist, being the first of its relata, as required for *this* respect of analogy with truth to go through. And I have said more about the R_k relation: I have called it the relation of vindication. So even if agreement with objective chance is just one acceptable answer to our puzzle, still we are committed to the existence of objective chances and a norm of veracity according to which credences strive to match them.

“Chance is mysterious”

Nevertheless, a number of authors find the very notion of chance to be incoherent (in the other sense of the word!), or at least mysterious. “Probability does not exist”, the great probabilist de Finetti quipped (1974, x)—meaning that *chance* does not exist, and that all probability is really subjective. Van Fraassen (1989) similarly finds objectionable an objective notion of chance—the sort of metaphysical spookery that should not have outlived the middle ages.

¹⁴ I owe this point to Paul Bartha.

I reply: Against the charge that chance is mysterious, one has to ask: *compared to what?* Probability is a modal concept, and as such is kindred to necessity, possibility, counterfactuals, and dispositions. All of these make an empiricist queasy, although they will not go away so easily. Even heirs to the empiricists such as de Finetti and van Fraassen need to populate their ontology with things that Hume would have no truck with. There is plenty to make Hume queasy, and not just Hume for that matter, in the notion of an infinite sequence of hypothetical trials. I am not sure that potential calibration has the empiricist high ground here.

In any case, I have given a new argument for believing that chance *does* exist. I now want to show that it need not be so mysterious, need not be so spooky.

For starters, here is an argument that chances must obey the probability calculus—and at least for this, we should thank calibration.¹⁵ It is a platitude, or close to it, that it is *possible* for chances to be perfectly calibrated (at least in the limit). Frequentists about chance think that it is not only possible, but analytic. We need not go that far. It suffices that frequencies *can* agree with chances. Then since chances do not ruin, a priori, the possibility of potential perfect calibration, they must satisfy the probability calculus. This is just as well, for credences strive to match them, and we have independent arguments that rational credences satisfy the probability calculus.

Or we could appeal to those independent arguments to reach the conclusion that chances must be probabilities by a different route. Suppose for reductio that chances do not obey the rules of probability, but rather obey some other rules, or no rules at all. Then credences must obey these other rules, or no rules at all, in order to match the corresponding chances. But in so doing, they disobey the probability calculus, and are *ipso facto* irrational (by those independent arguments). Thus, credences that have the virtue analogous to truth are *ipso facto* irrational. This is absurd—for all-or-nothing

¹⁵ Here I am especially grateful for discussion with Ned Hall.

beliefs that are true are not *ipso facto* irrational. Thus, we must reject our supposition, and conclude that chances *do* obey the rules of probability.

Or if you prefer (being already convinced of this conclusion), we can go in the other direction, giving a new argument for why credences should be probabilistically coherent: if they are not, necessarily some of them suffer from the probabilistic analogue of falsehood. This is analogous to the familiar argument that beliefs should be consistent: if they are not, necessarily some of them are false. The analogy provides yet more support for my claim that accordance with chances plays the role for credences that truth plays for beliefs: much as the norm of veracity imposes a structure on beliefs, codified by deductive logic, it imposes a structure on degrees of belief, codified by the probability calculus. And whichever way we run the argument, we have a welcome harmony between credences and chances, both of which are probabilities.

However, knowing that chances are probabilities does not yet tell us anything about their values, and it's their values that vindicate certain credences and vitiate others. Can we sometimes know, or at least be confident of, their values? I believe we can.

6. A probabilistic 'magic trick'

To show this, I want you to imagine my performing a 'magic trick'. I present you with a challenge:

1. Give me any object you like.
2. Give me any number between 0 and 1 inclusive (to be thought of as any chance value you like).
3. Give me any (finite) degree of approximation—any number of decimal places of accuracy you like.

My trick will be to use the object to generate an event whose chance is the number you chose, to the accuracy you specified. Your challenge is to stump me. I claim that you can't, modulo certain qualifications that I postpone until later, since I want to simplify my

initial exposition as much as possible. So let me put the point this cautiously: it is surprisingly hard for you to stump me.

Before giving the general procedure, and before confessing to the trick's limitations, let's work through one example, to see how the trick works.

1. You give me your car-key.
2. You choose the number $1/\sqrt{2}$.
3. You demand accuracy to three decimal places.

That is, you challenge me to produce an event whose chance is $1/\sqrt{2}$ at a 0.001 level of precision, using just your car-key.

I hope that at this stage you find it *surprising* that I can meet your challenge. For you have not made my task especially easy: I don't have any inside knowledge about your key, your choice of probability was not especially friendly, and you have demanded considerable precision.

Here's how the trick will go. 1) I will use the key to simulate a 'fair coin'; then, 2) I will use that 'fair coin' to generate the probability to the specified accuracy.

1) First, I look for some feature of the key that creates an asymmetry. For definiteness, I will suppose that there is a scratch on one side but not the other. Now, I toss the key a number of times, and record for each toss whether the key lands 'scratch UP' or 'scratch DOWN'. (Soon we will get a sense of how many tosses are required.) I produce a sequence of results—e.g.:

UP, DOWN, DOWN, DOWN, UP, UP, DOWN, UP, UP, UP, UP, DOWN, DOWN,
UP, DOWN, UP, DOWN, UP...

Now I divide the results into successive pairs:

<UP, DOWN>, <DOWN, DOWN>, <UP, UP>, <DOWN, UP>, <UP, UP>, <UP,
DOWN>, <DOWN, UP>, <DOWN, UP>, <DOWN, UP>, ...

Wherever I see the pair <UP, DOWN>, I write 'HEADS'; wherever I see the pair <DOWN, UP>, I write 'TAILS'. Wherever I see the other two pairs, I write nothing—

that is, I effectively discard the pairs <UP, UP> and <DOWN, DOWN>. In the example, I write:

HEADS, TAILS, HEADS, TAILS, TAILS, TAILS, ...

What is crucial here is that the *chance of heads equals the chance of tails*. After all, <UP, DOWN> has the same chance as <DOWN, UP>, irrespective of what the chance of UP is. (I'll say more about this shortly.) And given my rule on discarding, these are the only possible outcomes, so their chances conditional on having survived the discarding process are 1/2 each. So I have a sequence of trials, with a chance of 1/2 of each outcome on each trial. That is, I have *a simulation of a sequence of tosses of a fair coin*.

2) From now on, I will speak of the fair coin that I am simulating, rather than the key. The next step is to show that I can use the fair coin to generate the desired chance to the desired accuracy. You asked for $1/\sqrt{2}$ to three decimal places. That's 0.707, or 707/1000, to within 1/1000.

Note that $2^{10} = 1024 > 1000$. There are 1024 possible sequences of results of tossing a fair coin 10 times, and (given the coin's fairness), they are equiprobable. I can enumerate them in a long list:

HHHHHHHHHH

HHHHHHHHHT

HHHHHHHHTH

...

TTTTTTTTTT

Discard the bottom 24 sequences on the list; that leaves 1000 sequences. Draw a line immediately after the 707th sequence. Call any sequence above the line 'SUCCESS', below the line 'FAILURE'.

I now 'toss' the simulated fair coin 10 times. If the sequence of results happens to be one of the discarded 24 sequences, I rerun the experiment until I get either a 'SUCCESS' sequence, or a 'FAILURE' sequence. The chance of 'SUCCESS' is 707/1000—or more

precisely, the conditional chance of ‘SUCCESS’, given ‘SUCCESS’ OR ‘FAILURE’, is $707/1000$. That's $1/\sqrt{2}$ to three decimal places, and I used nothing but your car-key.¹⁶

We can now generalize the example, so that the trick works for *any* object, *any* chance, and *any* degree of approximation (within limits to be conceded shortly). As before, proceed in two stages: first, use the object to simulate a fair coin; then, use the fair coin simulation to generate the desired chance to the desired accuracy. The trick will work, at least in principle,¹⁷ if:

- i) there is some asymmetry that distinguishes two sorts of results that can occur;
- ii) I can generate a sufficiently long sequence of trials;
- iii) the chance of one sort of result remains constant from trial to trial;
- iv) the results of the trials are independent.

That is, it is guaranteed to work if I can generate a sufficiently long sequence of Bernoulli trials in which the chance of each of two possible outcomes is positive. In fact, I can replace the conjunction of iii) and iv) with a weaker assumption:

- iii)' the trials are exchangeable.

Trials are exchangeable (with respect to a probability function) if permuting them makes no difference to the probabilities of the outcomes—that is, the probabilities are insensitive to the order of the outcomes. If the trials are Bernoulli (with non-extreme probability), then they are exchangeable; however, the converse is false.

¹⁶ Alternatively, I could keep all 1024 sequences at the last stage, and draw a line immediately after the 724th sequence, since 724 is the integer closest to $1024/\sqrt{2}$. Call any sequence above the line ‘SUCCESS’, below the line ‘FAILURE’. Collectively, the SUCCESS sequences accrue $724/1024$ of the probability—that is $1/\sqrt{2}$ to three decimal places.

Recalling the observation in §4 that chances are really conditional on context propositions, it would be more accurate for me to say that the conditional chance of SUCCESS, *given* the key-tossing set-up, is $1/\sqrt{2}$ to three decimal places.

¹⁷ I add this hedge because there are ways in which you could thwart me in practice without undermining the spirit of the trick. For example, you might demand so many decimal places of accuracy that we would die before the experiment ended. Still, I can specify how the experiment would be performed.

It would be a little more accurate to say that the trick will *very probably* work if the conditions that follow are met. For I could be unlucky, and the object could yield the same result on every trial even though that result is chancy—the key *could* land ‘UP’ every time, however many times I toss it, even if the chance of ‘DOWN’ is positive on each toss. Nevertheless, the chance of my run of bad luck continuing diminishes, and in the limit vanishes, as the number of trials increases.

In fact, I do not even need to assume the full strength of exchangeability, since all I need is invariance of the probabilities under permutations of pairs of trials. I don't believe this property has a name, so let me give it one: *pairwise exchangeability*. So my trick is guaranteed to work if as well as i) and ii),

iii)" the trials are pairwise exchangeable.

(I will say more about assumptions iii)' and iii)" shortly.)

At this point, I am using the given object to simulate the results of tossing a coin of unknown, but fixed, bias. The pairing and discarding process turns a 'biased coin' into a 'fair coin'.¹⁸

Having now a simulation of the tosses of a fair coin, the second stage is to use this to generate the desired chance to the desired accuracy. I can determine how long a sequence of tosses will produce the desired accuracy. If you want accuracy to 6 decimal places, 20 tosses suffice, since $2^{20} > 10^6$. In general, if you want accuracy to n decimal places,

$x =$ the smallest integer greater than $n \cdot \log_2(10)$

tosses suffice, since $2^x > 10^n$. List all sequences of x tosses; discard all sequences that fall below the (10^n) th place; express your desired chance as a decimal d (to n places); draw a line after the $(d \cdot 10^n)$ th place, and call any sequence that falls above the line a 'SUCCESS'. The chance of a SUCCESS is what you desired, to the desired accuracy, and I used nothing but your object.

Now, I do not claim to have created chances *ex nihilo*. That would be *quite* a trick! Rather, my trick takes us from 'coarse-grained', qualitative inputs to 'fine-grained', quantitative outputs—somewhat reminiscent of the way Ramsey (1926)/Savage (1954)/Jeffrey (1966) take us from qualitative inputs concerning preferences to quantitative utility and probability functions. My trick piggy-backs on chances that are

¹⁸ I learned from Paul Vitanyi that von Neumann showed that one can use any given coin to simulate a fair coin. I thank Branden Fitelson for giving me the reference: von Neumann (1951).

already out there when I assume that the trials are pairwise exchangeable. I admit that this assumption is non-trivial, and I will soon countenance various systems for which this or one of the other assumptions fail. Still, I maintain that for a surprisingly large range of objects, I can make the assumptions hold and the trick work. And once you agree that it works, there is little wiggle room: the probability that you should assign to the event in question is tightly constrained. While I don't show that every experiment has an associated chance, I do claim to show that every chance has an associated experiment (up to the accuracy required).

In any case, physics and commonsense apparently tell us that various events *are* independent of each other and that their probabilities remain fixed, or very nearly fixed, from trial to trial. Still more of them are exchangeable. Still more of them are pairwise exchangeable. One can always feign skepticism about these facts, as one can feign skepticism about just about any fact; we've all read Descartes, Hume, and Goodman. But typically such skepticism is *only* feigned. A sensible Bayesian will, I hope, agree with me that certain probability assignments—for example, $1/\sqrt{2}$ to three decimal places for the event I constructed—are natural, and possibly even required given our shared background knowledge of the way the world works. If *you* are skeptical of that, please contact me. Let's make some bets. Don't forget to bring your car-key.

7. Confessions, concessions, and comebacks

Even if you remain skeptical about my magic trick in most cases, as long as you allow me that it *sometimes* works, my job is done. I have been arguing, remember, that non-trivial chances need not be mysterious—that we can *sometimes* have a good handle on their values. Even if these cases are rare, their existence is all I need. For once we are dealing with such a case, the Bayesian cannot blithely say that all coherent degrees of belief are equally 'true'. Suppose, for definiteness, that the car-key example is such a

case. Then if I assign probability 0.707 to the event I constructed, and you assign 0.123 to it, we are disagreeing. More than that, to put it bluntly, *I am right and you are wrong*.

Now I'm sure that if you set your mind to it, you could think of ways of ensuring that my assumptions fail in a given case, thus stymieing me. Here are some ways.

Systems that lack a discernable asymmetry

A system that lacks any discernable asymmetry might thwart me at the very outset—for example, a sphere on which no blemish can be detected. For then I may not be able even to generate two kinds of outcomes that I can discern. Now, I might be able to use the immediate environment to break the symmetry—for instance, by rolling the sphere on a tiled floor, and observing whether it stops to the left or to the right of a given line. You might thwart me again by making the environment symmetric too—say, by placing the sphere inside a featureless rocket outside a detectable gravitational field... But this is getting far-fetched; you will be hard-pressed to hand me such a system in the real world. And we would have no use for probability for such a system in any case, so I will consider it no further.

Systems for which I cannot generate multiple trials

More realistically: fragile or short-lived systems will ruin my trick because I will not be able to generate the run of trials that I require. Indeed, the notion of pairwise exchangeability does not even make sense for a very fragile or very short-lived object that only survives a single trial.

Systems with memories

Then there are systems for which we can identify two kinds of outcomes on repeated trials, so that pairwise exchangeability makes sense—but fails. You might, for example, hand me your cat as the object, and defy me to use her to generate a sequence of pairwise

exchangeable trials. Or consider someone learning to throw a dart at a bull's eye, who can either HIT or MISS it on a given throw: the sequence <MISS, HIT> could be a bit more probable than <HIT, MISS>, because the dart-thrower's accuracy improves slightly with practice. So the trials are not pairwise exchangeable. When a system has a memory and modifies its behavior on the basis of what has happened in the past—as the cat and the dart-thrower do—it may foil the assumption of pairwise exchangeability.

A world with extreme chances, and enclaves with extreme chances

Suppose, as we have before, that all chances in the world are 0 or 1. Then *all* objects will fail to yield pairwise exchangeable trials unless they yield the same outcome every time, in which case my trick is frustrated in any case. And if there are enclaves in which all chances are 0 or 1, then objects selected from these enclaves will frustrate my trick. That's because pairwise exchangeability fails. Suppose that all chances concerning my key-tossing are 0 or 1, and that in fact the first toss lands UP, and the second lands DOWN. These events are not pairwise exchangeable, because <UP on first toss, DOWN on second toss> has chance 1, while <DOWN on first toss, UP on second toss> has chance 0.

Almost pairwise exchangeability—but not quite

Or you might grant me the intermediate chances that I need, and grant me that the relevant trials are *almost* pairwise exchangeable, but not quite. You say that the key landing <UP, DOWN> has *almost* the same probability as <DOWN, UP>, but not quite—the two might differ in the fourth decimal place, say. (The key is like a dart-tosser who improves very slowly, or a drinking dart-tosser who gets worse very slowly!) So the coin that I will simulate will be *almost* fair, but not quite. And the small deviation from fairness will foil my claim to deliver chances to any precision you want. At best, I will only be able to deliver chances to within intervals of the target values—and the further

the relevant trials are from being perfectly pairwise exchangeable, the wider the intervals will be.

I may be able to refine my putative 'fair coin' for you, and thus improve my precision. Suppose you are skeptical that I have really created a 'fair coin' out of the car key: <UP, DOWN> does not have *exactly* the same probability as <DOWN, UP>, you say. Still, you may well agree that their probabilities (conditional on having survived the discarding process) are very close to 1/2, and in particular much closer than the probabilities of UP or of DOWN were. Then simply run the trick again, driving the probabilities still closer to 1/2. Now call 'HEADS' the pair <<UP, DOWN>, <DOWN, UP>>, and 'TAILS' the pair in which the ordered pairs are reversed: <<DOWN, UP>, <UP, DOWN>>. This should be a still fairer 'coin' than the one I originally created. And so on, making ever fairer the 'coin', as need be. More generally, you may well agree that my trick often drives probabilities much closer to 1/2, and that by repeatedly rerunning the trick, we can drive them ever closer again.

Still, this procedure has its limits, and there are doubtless systems for which it will not work at all (I have already considered some). Very well, then; the trick is not foolproof—or catproof, or dart-thrower-proof, or To that extent, I cannot make a bold, unqualified promise to fulfill the magic trick come what may (not that I ever did). But far from weakening my overall position, I insist that it only reinforces my main points.

Firstly, bad news for my magic trick is more bad news for frequentism, and thus for the friend of calibration. If you do *not* grant my assumptions regarding a particular case, you should not be a frequentist or a calibrationist regarding that case, either. If the object is fragile or short-lived, then the number of trials will be minimal: relative frequencies are trivialized, and calibration is nugatory. (At the extreme, we have the problem of the single case for frequentism and for calibration.) And if the trials are not Bernoulli, nor even exchangeable, nor even pairwise exchangeable, then frequencies are a poor guide to

the probability on a *given* trial. Frequencies are *insensitive* to trial order; according to the relative frequency function, the trials are exchangeable, and *a fortiori* pairwise exchangeable. But when the trials are not exchangeable, nor even pairwise exchangeable, the probabilities are *sensitive* to their order. Suppose, for example, that our dart-thrower learns from previous trials, so that <MISS, HIT> is more likely than <HIT, MISS>. But the frequentist who equates the probability of ‘HIT’ with its relative frequency in some run of trials accords them the same probability. And the probability of ‘HIT’ on the first trial could obviously come apart from its relative frequency in the run of trials. This, in turn, bodes ill for the cogency of calibration in such cases.¹⁹

Note that chances are not susceptible to these worries, nor is my proposal that they are the vindicators of credences. Chances are not held hostage to there being multiple trials. Trials may or may not be pairwise exchangeable according to the chance function, and the ‘true’ credences will yield the same verdict. Indeed, the whole point of my latter concessions above was that some systems might thwart my magic trick because the relevant chances are not pairwise exchangeable. But in raising these objections to the assumptions underlying the magic trick, I have really been setting a trap for the subjective Bayesian. By all means, protest—the shriller, the better!—that my exchangeability assumption fails for some object that you hand me (a cat, a dart-thrower, or what have you). For if you do so, you are really my ally, and you are no friend of the orthodox Bayesian.

That is because, secondly, bad news for my magic trick is more bad news for the orthodox Bayesian. What does it mean for my exchangeability assumption to *fail*? Pairwise exchangeability is a *relation* between a set of trials and a probability function. Fix the set of trials, and the relation will hold for some probability functions and not others. If you merely say that the trials generated by some object are not pairwise

¹⁹ See Seidenfeld (1985) for a nice discussion of how calibration is an unreasonable goal when the trials are not independent.

exchangeable *according to your subjective probability function*, then I thank you for that piece of autobiography, but I am not yet troubled. Instead, I have imagined you claiming that the trials generated by the cat, the dart-thrower, or what have you, are not pairwise exchangeable *in fact*. The failure of that exchangeability is supposed to be *out there in the world, irrespective of what I think*. If I disagree with you about the holding of the exchangeability assumption, you think that *you are right and I am wrong*. However, these are not things that the orthodox Bayesian can say. By all means, declare that it is *an objective fact* that the pairwise exchangeability assumption fails. But then it sounds as if you have in mind pairwise exchangeability relative to *the objective probability function*. I prefer to call it the chance function.

8. Conclusion

I do not claim to have given an analysis of chance—that is avowedly a topic for another occasion. I *do* claim to have given a new argument for believing that chances exist, and moreover to have shown that we can be confident of the chances of at least some propositions. And I claim to have solved the puzzle of this paper: a degree of belief is ‘true’—vindicated—when it agrees with the corresponding chance (rather than when it is part of a perfectly calibrated set of credences). I contend that Bayesianism should heed all of this. Now, you may think that absent an analysis of chance, we have no way of assessing my claims. I demur. I have no analysis of *truth* handy, and you may not have one handy either. But even in these post-modern times we can agree that truth exists, that we can be confident of the truth-values of at least some propositions, and that beliefs are vindicated by it. Why hold chance to a higher standard?²⁰

²⁰ I am grateful for very helpful comments from Phillip Catton, Mark Colyvan, Andy Egan, David Lewis, Peter Menzies, Ralph Miles, Andrew Reisner, Elliott Sober, Manuel Vargas, Susan Vineberg, Paul Vitanyi, Brian Weatherston, Lyle Zynda, and especially Paul Bartha, Alex Byrne, Fiona Cowie, David Dowe, Adam

REFERENCES

- Belnap, Nuel and Mitchell Green (1994), “Indeterminism and the Thin Red Line”, *Philosophical Perspectives 8: Logic and Language*, ed. James Tomberlin, Atascadero: Ridgeview, 365-88; revised as chapter 6 of Belnap et al., *Facing the Future*.
- Belnap, Nuel, Michael Perloff and Ming Xu (2001), *Facing the Future: Agents and Choices in Our Indeterministic World*, Oxford: Oxford University Press.
- Cartwright, Nancy (1999), *The Dappled World: A Study of the Boundaries of Science*, Cambridge: Cambridge University Press.
- Churchland, P. M. (1981), “Eliminative Materialism and the Propositional Attitudes”, *The Journal of Philosophy* 78: 67-90.
- Dawid, A. P. (1982), “The Well-Calibrated Bayesian”, *Journal of the American Statistical Association*, 77 (379): 605-610.
- de Finetti, Bruno (1974), *Theory of Probability*, Vol. I, Wiley, New York.
- Earman, John (1992), *Bayes or Bust?*, MIT Press.
- Giere, Ronald N. (1973), “Objective Single-Case Probabilities and the Foundations of Statistics.” In *Logic, Methodology and Philosophy of Science IV*, edited by P. Suppes, et al., New York: North-Holland, 467-483.

Elga, Branden Fitelson, Ned Hall, Matthias Hild, Christopher Hitchcock, Marc Lange, Brad Monton, Daniel Nolan, Michael Thau, Peter Vranas, Jim Woodward, and Ralph Wedgwood. Early versions of this paper were given as talks at the University of Melbourne, Cambridge University, Columbia University, the San Diego meeting of INFORMS, the Australasian Association for the History, Philosophy, and Social Studies of Science conference in Melbourne, and the Australasian Association of Philosophy conference in

- Hájek, Alan (2003a), “Conditional Probability is the Guide to Life”, in *Probability is the Very Guide of Life: The Philosophical Uses of Chance*, eds. Henry Kyburg, Jr. and Mariam Thalos, Open Court, 183-203. Abridged version in *Proceedings of the International Society for Bayesian Analysis 2002*.
- Hájek, Alan (2003b), “What Conditional Probability Could Not Be”, *Synthese* 137 (3).
- Harman, Gilbert (1986), *Change in View: Principles of Reasoning*, MIT Press.
- Jeffrey, Richard (1966), *The Logic of Decision*, Chicago: University of Chicago Press; 2nd ed. 1983.
- Jeffrey, Richard (1977), “Mises Redux”, in R. E. Butts and J. Hintikka (eds.), *Basic Problems in Methodology and Linguistics*, Reidel, 213-222. Reprinted in *Probability and the Art of Judgment*, Cambridge: Cambridge University Press, 1983, 192-202.
- Kyburg, Henry E. (1968), “Bets and Beliefs”, *American Philosophical Quarterly*, 5 (1): 54-63.
- Lad, Frank (1996), *Review of Operational Subjective Statistical Methods: A Mathematical, Philosophical, and Historical Introduction*, New York: Wiley.
- Lange, Marc (1999), “Calibration and the Epistemological Role of Bayesian Conditionalization”, *The Journal of Philosophy* 96: 294-324.
- Levi, Isaac (1990), “Chance”, *Philosophical Topics*, vol. 18, No. 2: 117-148.
- Lewis, David (1980), “A Subjectivist's Guide to Objective Chance”, in *Studies in Inductive Logic and Probability*, Volume II., University of California Press, 263-

Adelaide, and the paper has profited from comments by various audience members—especially Otavio Bueno, Peter Forrest, David Lewis, Christian List, Graham Oppy, Denis Robinson, and Neil Thomason.

- 293; reprinted in *Philosophical Papers Volume II*, Oxford: Oxford University Press, 1986, 83-132.
- Lewis, David (1994), "Humean Supervenience Debugged", *Mind*, 103: 473-490.
- Popper, Karl (1959), "The Propensity Interpretation of Probability", *British Journal of the Philosophy of Science* 10: 25-42
- Ramsey, F. P. (1926), "Truth and Probability", in *Foundations of Mathematics and other Essays*, R. B. Braithwaite (ed.), Routledge & P. Kegan , 1931, 156-198; reprinted in *Studies in Subjective Probability*, H. E. Kyburg, Jr. and H. E. Smokler (eds.), 2nd ed., R. E. Krieger Publishing Company, 1980, 23-52; reprinted in *Philosophical Papers*, D. H. Mellor (ed.) Cambridge:University Press, Cambridge, 1990.
- MacFarlane, John (2003), "Future Contingents and Relative Truth", forthcoming in the *Philosophical Quarterly*.
- Mellor, Hugh (1971), *The Matter of Chance*, Cambridge: Cambridge University Press.
- Savage, Leonard J. (1954), *The Foundations of Statistics*. New York: John Wiley.
- Seidenfeld, Teddy (1985), "Calibration, Coherence, and Scoring Rules", *Philosophy of Science* 52: 274-294.
- Shimony, Abner, (1988), "An Adamite Derivation of the Calculus of Probability", in J. H. Fetzer (ed.), *Probability and Causality*, Dordrecht: D. Reidel, 79-89.
- Skyrms, Brian (1984), *Causal Necessity*, New Haven: Yale University Press.
- van Fraassen, Bas (1983), "Calibration: A Frequency Justification for Personal Probability", in R. S.Cohen and L. Laudan (eds.), *Physics, Philosophy and Psychoanalysis*, Dordrecht: D. Reidel Publishing Company, 295-319.

van Fraassen, Bas (1984), "Belief and the Will", *Journal of Philosophy* 81: 235-256.

van Fraassen, Bas (1989), *Laws and Symmetry*, Oxford: Clarendon Press.

von Mises, Richard (1957), *Probability, Statistics and Truth*, revised English edition,
New York: Macmillan.

von Neumann, J. (1951), "Various Techniques Used in Connection with Random Digits",
Monte Carlo Method, *Applied Mathematics Series*, No. 12, U.S. National Bureau
of Standards, Washington D.C. 36-38.