

Rationality and Indeterminate Probabilities

Alan Hájek and Michael Smithson¹

Australian National University

Introduction

What is your degree of belief that the Democrats will win the next presidential election in the USA? If you report a sharp number, we will question you further. For example, if you report a credence of 0.6, we will ask whether you really mean 0.6000..., sharp to infinitely many decimal places. If you are anything like us, your credence is sharp only up to one or two decimal places. And in that case, you are not an ideal Bayesian agent. For such an agent assigns perfectly sharp credences to all propositions.

Some Bayesians may nevertheless let you into their fold, more or less grudgingly. For they are prepared to countenance *indeterminate* credences—just how grudgingly will vary from one Bayesian to another. Jeffrey (1992), for example, is quite prepared to give Bayesianism “a human face”, and he believes that human credences may be more realistically modelled as a set of probability functions. Still, grudging or not, it seems to be something of a concession: his preferred model of an ideally rational agent portrays her as having a single probability function.

In this paper, we will argue that indeterminate probabilities are not only rationally permissible, but they may even be rationally *required*. We follow Levi’s (2000)

¹ We thank especially Rachael Briggs, Mark Colyvan, John Cusbert, Kenny Easwaran, Adam Elga, Christoph Fehige, Aidan Lyon, Elijah Millgram, Susanna Rinard, Jonathan Schaffer, Teddy Seidenfeld, Nick Smith, Katie Steele, Larry Temkin, Roger White, and the audience at an RSSS seminar at the Australian National University for very helpful discussion. Special thanks to Aidan Lyon and Wolfgang Schwarz for detailed and penetrating comments on an earlier draft. We also thank Elle Benjamin and Ralph Miles for editorial assistance.

important distinction between the *elicitation* of an agent's credal state, and the state *itself*, and we follow him again in using the word “indeterminate” to describe a feature of the state itself—its consisting of more than one probability function, as we will say (for example, upper and lower probability functions, or a non-singleton set of probability functions).

There are already a number of arguments in the literature for being receptive to indeterminate credences. These include:

- Mathematical considerations (see e.g. Walley 1991: 67-86, 313-317, 317-322, and 328-330; and Seidenfeld and Wasserman 1993: 1139).
- Psychological considerations, especially in light of preference patterns that Ellsberg (1961) identified (see e.g. Smithson 1999).
- Partition-independence of prior probabilities (Walley 1991, 227-8; 1996).
- The handling of group decision problems (see e.g. Levi 1982, Seidenfeld et al. 1989).

Upon reflection on all these considerations, we are convinced that the case for recognizing indeterminate credences is already good. We hope now to make it even better.

Our first argument begins by assuming a version of *interpretivism*, an influential position in the philosophy of mind: your mental state is the set of probability and utility functions that rationalize your behavioral dispositions as well as possible. But consistent with your being rational, the set of best-rationalizing functions may consist of multiple probability functions—there is no further fact that would single out any of them. Then according to interpretivism, this *makes it the case* that your credal state is indeterminate.

The second argument is based on the possibility of indeterminate *chances*. We describe a world that plausibly has indeterminate chances. Moreover, a version of a chance-credence coordination principle, well-known to philosophers as Lewis's (1980) 'Principal Principle', requires a certain alignment of a rational agent's credences with corresponding hypotheses about the chances. Thus, if the chances are hypothesized to be indeterminate, the agent will inherit their indeterminacy in her corresponding credences.

The third and longest argument is motivated by a dilemma. Epistemic rationality requires you, among other things, not to be *dogmatic*—to stay open-minded about contingent matters about which your evidence has not definitively legislated. Practical rationality requires you, among other things, not to be *paralyzed*—to be able to act decisively at least sometimes. It turns out that these requirements can conflict with each other, if we model open-mindedness as assignments of positive credences, and rational decision as maximization of expected utility. For thanks to your open-mindedness, some of your options may have *undefined* expected utility, and if you are choosing among them, decision theory has no advice to give you. One such option is playing the *Pasadena Game*, a St. Petersburg-like game introduced by Nover and Hájek (2004; see also Hájek and Nover 2006, 2008), which we will rehearse. Moreover, any option that yields *positive probability* of playing the Pasadena Game is also undefined in expected utility; but as we will argue, this may well be true of all your options thanks to your open-mindedness—in which case by the lights of decision theory, you are paralyzed.

Indeterminate probabilities to the rescue! A way out of your predicament is to serve both masters, epistemic rationality and practical rationality, in one fell swoop with an indeterminate credence to the prospect of playing the Pasadena game. You

serve epistemic rationality by making your upper probability positive—it ensures that you are open-minded. You serve practical rationality by making your lower probability 0—it provides guidance to your decision-making. No sharp credence could do both; we thus have a new kind of argument for indeterminate credences.

Sharpening up our terminology

The terminology in this area is a philosophical minefield, so we need to sharpen it up. Various words of related but distinct meanings have been used: *indeterminate*, *imprecise*, *vague*, *indefinite*, *fuzzy*, *mushy*, ... Moreover, these words of ordinary English also have proprietary technical meanings. For example, the word ‘vague’ has philosophical connotations that would be misleading here. Vagueness is typically characterized by sorites-susceptibility, and by the existence of borderline cases whose classification is somehow problematic. Moreover, many philosophers think that where there is vagueness there is necessarily higher-order vagueness: there is vagueness regarding what the borderline cases are. The words ‘indefinite’, ‘fuzzy’ and ‘mushy’ also suggest higher-order vagueness. But none of these phenomena is displayed by a credence of $[0.5, 0.7]$, and yet it is a paradigm case of the sort of credence we want to characterize.

Up to a point, it is a matter of convention which terminology we adopt. But in doing so we must be careful not to collapse important distinctions. Above all, we should distinguish lack of sharpness in our *elicitation* or measurement of an agent’s credences, with lack of sharpness *in the credences themselves*. Levi (2000) is quite clear on this distinction. What he calls “imprecise” probabilities arise from difficulties in elicitation: an agent has a unique subjective probability function, but she (or another ascriber) cannot figure out exactly what it is. Her credal state is in fact

perfectly sharp, but there is some epistemic obstacle to accessing it; she (or the ascriber) simply doesn't know her mind. (This is reminiscent of “epistemicism” about vagueness, as defended by Williamson 1994 and Sorensen 2001—the view according to which, for example, there is an exact point at which adding a grain of sand turns a non-heap into a heap, but we are ignorant what it is.) Levi’s terminology accords well with the familiar usage according to which “imprecision” is usually thought of as a property of measurement. What Levi calls *indeterminate probabilities*, by contrast, arise when the agent’s credal state is *itself* not sharp. She (or the ascriber) may know her mind perfectly well, and the indeterminacy of probability assignments resides there. Walley’s magnificent book (1991), for example, conflates these notions: some of his arguments for what he calls “imprecision” concern an agent’s mental state itself, while others concern its elicitation. Perhaps partly due to the influence of Walley, the *Society for Imprecise Probability: Theory and Applications* has even adopted the word ‘imprecise’ in its very name! And yet its purview again includes both imprecision and indeterminacy.

So when we say that a credal state is indeterminate, we are denying that it corresponds to a single probability function that takes numbers as values. We remain neutral as to how exactly such a credal state should be represented. There are various mathematical representations of indeterminate credences—as lower and upper probabilities, as intervals, as convex sets of probability functions, as possibly sparser sets of such functions, etc. Interesting though intra-mural disputes among proponents of these representations are, we will not enter into the fray. We are happy to remain ecumenical about these approaches; perhaps, for example, different sources of indeterminacy motivate different representations. We do not need to take a sharp stand here on the nature of indeterminacy. It will be enough for our purposes if we

can convince you that *some* such representation is required in our theory of rational credence. That said, it will sometimes be convenient for us to speak in terms of upper and lower probabilities, or intervals, or sets of probability functions, to make a point vivid.

Interpretivism about mental states

A central issue in the philosophy of mind concerns what it is for an agent to possess a particular mental state. *Interpretivists* identify the state with the best possible *judgment* of the state by an ideal interpreter. For example, according to Lewis (1974), your mental state *is* the probability function/utility function pair that best rationalizes your behavioral dispositions—rationalization understood as expected utility maximization. We may imagine an ideal interpreter fitting various such pairs to you, and determining which fit best; but fixating too much on this image might mislead one into thinking that our topic is elicitation, when it is not. Rather, this is an account of what the mental state *is*. Now, typically there is a limit to how fine-grained these dispositions can be, allowing a multiplicity of equally good interpretations—thus, we argue, rendering your credences indeterminate by interpretivist lights. More generally, rationality surely permits an agent to have dispositions that admit of multiple equally good interpretations; but by the lights of interpretivism, this means that rationality permits indeterminate credences.

Note well: this is *not* to conflate imprecision (in Levi's sense) and indeterminacy. For example, Jeffrey (1992) does appear to conflate these two notions, when arguing that when your preferences do not determine a unique probability function (and utility function), your credences are indeterminate (in the terminology that we are using). Interpretivism is essential to our argument, but plays no role in his. The crucial point

is that the distinction between imprecision and indeterminacy collapses for the *ideal* interpreter (although of course not in general). If this interpreter attributes multiple probability functions to the agent, then this *makes it the case* that her *credences themselves* are indeterminate. It's not that there is some further fact that determines a unique function from the set of best-interpreting functions, of which the interpreter is ignorant. All the relevant facts are in. None of the functions is privileged.

To be sure, you may run the argument in reverse: the fact that the distinction between imprecision and indeterminacy collapses for the ideal interpreter is a *reductio* of interpretivism, you may say. In that case, you may still agree with the conditional: *if* interpretivism is correct, then rationality permits indeterminate credences. And interpretivism has its share of advocates among philosophers of mind; they should apply modus ponens and discharge the consequent.

Chances and credences

Lewis's interpretivism about credences resonates with his analysis of objective *chances* (1994). He identifies chances with their role in the best theory of the universe—one that best combines simplicity, strength, and fit to the data. Thus, much as he analyzes your credences as those probabilities attributed to you by the ideal interpreter, we may think of the chances as those probabilities attributed to the universe by the ideal science—the best interpreter of the universe, so to speak.

Lewis observes that a theory might gain greatly in simplicity, without much loss in strength, by positing stable chances over outcomes, rather than stating the outcomes themselves. For example, consider two competing theories of a universe consisting of a single coin that is repeatedly tossed. The first theory lists each outcome, toss by toss. It is a strong but highly complicated theory. The second theory states that

‘Heads’ and ‘Tails’ each have a constant chance of $1/2$. It is a weaker but much simpler theory. Lewis’s point, as illustrated by this example, is that the second theory might on balance be the superior theory, and indeed might be the best theory of this universe. Then *this makes it the case* that the universe’s laws are indeterministic, and the chances are just what the theory says they are.

Now suppose that the relative frequencies for some event-type vary erratically over time, but stay confined to a certain interval—say, $[0.4, 0.6]$. We can imagine this pattern persisting forever, so that there is no limiting relative frequency, although asymptotically the relative frequencies stay in this interval. Rather than positing time-dependent chances, the best system might posit a stable indeterminate chance of $[0.4, 0.6]$ for this event-type.² But then on the Lewisian picture, that *makes it the case* that the true chance is indeterminate over this interval. And whether or not this is actually the case, a rational agent can conditionally entertain the proposition that it is. It is unclear whether there are real-world examples of such phenomena. Perhaps Papamarcou and Fine (1986, 711) have one: “flicker or $1/f$ noise provided by the frequency fluctuations of high quality quartz crystal oscillators used in conjunction with atomic clocks”—although to be sure, such frequencies are finite, so it is hard to assess their limiting behavior. But whether or not the real world furnishes such examples, we should be receptive to them in our theorizing. After all, nobody claims that the St. Petersburg game is found in the real world, yet it has been a touchstone for much work in decision theory, and rightly so (and soon enough we will be invoking another St. Petersburg-like game).

² This is inspired by a remark in Walley and Fine (1982): “Just as “randomness” (chance) is introduced in additive probability models to account for poorly understood (“accidental”) variation in outcomes, so “imprecision” might be introduced in additive probability models to account for poorly understood variations in chance behaviour” (759). But our argument is not based on any epistemic considerations, as suggested by their words “poorly understood”.

In any case, it seems that rationality allows, and maybe even requires, the typical Bayesian agent to give some credence to the hypothesis that at least some chances are indeterminate. This is clearest for the prior probability assignment of a Bayesian agent who lacks any evidence that would conclusively rule out such a hypothesis; and it is plausible enough even for an agent who is well along her Bayesian odyssey, but whose accumulated evidence still does not conclusively rule out such a hypothesis. Now assume a version of a chance-credence coordination principle, such as Lewis's well-known Principal Principle (1980), according to which your credence in a proposition, conditional on the chance of that proposition being x , should be x . (See Lewis for further fine-tuning.) Extend this in a natural way to allow for indeterminate chances: your credence in a proposition, conditional on the chance of that proposition being indeterminate in a particular way, should be indeterminate in the same way. The indeterminacy in the chance (either actual, or entertained) that appears in the condition is inherited by your conditional credence.

We thus have an argument that ideal rationality permits, and maybe even requires, indeterminate credences. Ideal rationality permits, and maybe even requires, at least some agents to give positive credence to at least some indeterminate chance hypotheses. For definiteness, let one such hypothesis take the form

$$\text{chance}(X) = [p, q], \text{ where } 0 \leq p < q \leq 1,$$

and let you be the agent. Then by a natural extension of the Principal Principle, your credence in X , *conditional* on this hypothesis, is correspondingly indeterminate:

$$P(X \mid \text{chance}(X) = [p, q]) = [p, q].$$

If you regard the chance function as indeterminate regarding X , it would be odd, and arguably irrational, for your credence to be any sharper. Compare: if your doctor is your sole source of information about medical matters, and she assigns a credence of

[0.4, 0.6] to your getting lung cancer, then it would be odd, and arguably irrational, for you to assign this proposition a sharper credence—say, 0.5381. How would you defend that assignment? You could say “I don’t *have* to defend it—it just happens to be my credence.” But that seems about as unprincipled as looking at your sole source of information about the time, your digital clock, which tells that the time rounded off to the nearest minute is 4:03—and yet believing that the time is in fact 4:03 *and 36 seconds*. Granted, you may just happen to believe that; the point is that you have no business doing so.

Here is another way in which chances could be indeterminate. Much as there could be some leeway in the best interpretation of you, so there could be some leeway in the best interpretation of the universe. Now imagine two different theories tied for first place in the Lewisian competition for the best system. They have equal claim to dictating what the chances are, but they disagree on what they are. Then we might say that to this extent it is indeterminate what the chances are. There is still more opportunity for indeterminacy if more theories are tied. Note well: paralleling our discussion in the previous section, this is *not* to conflate the epistemological issue of *eliciting* what the true chances are, with the metaphysical issue of *what they are*. The crucial point is that the distinction between the elicitation of chances and what the chances are collapses for *ideal* science (although of course not in general). If multiple chance functions are attributed to the universe by equal-best theories, then this *makes it the case* that *the chances themselves* are indeterminate. And as before, you may run the argument in reverse: the fact that the distinction between ideal elicitation of the chances and the chances themselves collapses is a reductio of the Lewisian view of chance, you may say. In that case, you may still agree with the conditional: *if* his view is correct, then the chances may be indeterminate. And the Lewisian view has its share

of advocates among philosophers of science; they should apply modus ponens and discharge the consequent.

Be that as it may, we ask that you grant us for now the possibility that some particular chance is indeterminate—offhand this seems at least to be a coherent hypothesis (and if Papamarcou and Fine are on the right track, then that may be an understatement). As such, a rational agent may dignify it with some positive credence. Then she will inherit the hypothesized indeterminacy in her corresponding conditional credences via a natural extension of the Principal Principle.

* * * * *

Now that we have discussed how indeterminate probabilities can resolve conflicting rationality norms, we are in a position to provide our final, somewhat lengthy argument for indeterminacy. We do not claim that it is decisive—as always, there will be possible responses, and we will countenance some ourselves. Still, we offer it as a new kind of motivation for indeterminate credences, quite unlike any that we have seen before.

The Pasadena Game

Our final study will involve the apparent conflict of two further rationality norms:

- *Open-mindedness*: you should not assign (sharp) probability 0 to any possibility that your total evidence does not rule out.
- *Some Expected Utility Maximization*: at least some of your actions should maximize your expected utility.

We may defend the norm of open-mindedness by arguing that to violate it is to misrepresent your evidential state—to treat something compatible with your evidence as if it were incompatible. Moreover, any violation of open-mindedness involves

treating an doxastically live proposition as if it were as dead as a logical contradiction.³ Some probabilists advocate a stronger norm than open-mindedness, sometimes given the strikingly unevocative name “regularity”: you should not assign probability 0 to any possibility. “Keep the door open, or at least ajar”, advise Edwards, Lindman and Savage (1963) in their defense of regularity. This requires you to stay open-minded about a possibility even when you have decisive evidence against it. It would prohibit you from ever conditionalizing on some proposition—for once you do, your mind is closed regarding all possibilities incompatible with that proposition. Our strictly weaker open-mindedness requirement only prohibits you from closing your mind to possibilities that are consistent with all your evidence. If you gain evidence E , then you are free to conditionalize on it, *contra* regularity; to be sure, you then assign probability 0 to all $\neg E$ possibilities, but your total evidence rules them out. Open-mindedness merely prohibits leaps of unfaith.⁴

³ To be sure, your different attitude to a doxastically live proposition L and some contradiction C might be revealed in your *conditional* probabilities, and in your updating dispositions. For example, your conditional probability for L , *given* L , should be 1, whereas your conditional probability for C , *given* L , should (still) be 0. Note that this requires some modification to Bayesian orthodoxy, according to which conditional probabilities—as defined by the usual ratio formula—are not defined when the condition has probability 0. See Hájek (2003) for further discussion of this issue, and for a defense of primitive conditional probability functions—in particular, Popper functions.

While granting that your different attitudes to L and C might be revealed in your conditional probabilities, it is still troubling that they are not *also* revealed in your *unconditional* probabilities. After all, these represent your attitude to the world absolutely, not under any condition or supposition. Surely your *unconditional* view of the world distinguishes L and C ; it isn’t merely under some condition or supposition that you should have different degrees of confidence in them. L is doxastically live for you, while C is doxastically dead for you, *unconditionally*.

⁴ You may think that quick counterexamples to Open-mindedness are provided by non-empty sets of measure 0. For instance, if you are throwing a dart at a representation of the $[0, 1]$ interval, you may (and perhaps even should) give probability 0 to the dart hitting the point $\frac{1}{2}$, or indeed to its hitting a rational number. In reply, note that Open-mindedness may be saved here by allowing *infinitesimal* probability assignments. (See Bernstein and Wattenberg 1969.) In any case, whatever the putative technical difficulties standard probability theory may have with

The norm of Some Expected Utility Maximization may strike you as surprisingly coy—after all, we are usually told that rationality requires *all* of your actions to maximize expected utility. If it does, all the better for the plausibility of the norm; it does not impugn its truth that a stronger norm is also true. But it is not clear that the stronger norm is true. For starters, we may imagine cases in which you have infinitely many options, none maximal in expected utility. For example, you may get to choose how many days you get to stay in heaven, only finite numbers allowed! Or as we are about to see, we may imagine curious cases in which the expected utility for at least one of your options is apparently undefined, so that none of your choices maximizes your expected utility. Such considerations may caution us not to endorse the stronger norm unrestrictedly. And so we do not. The weaker norm will be problematic enough.

The rationality norms are before you; now let's see how they apparently collide. This will involve a curious gamble called the *Pasadena game* (Nover and Hájek 2004, Hájek and Nover 2006, Hájek and Nover 2008). It resembles the St. Petersburg game in two important respects: a fair coin is tossed until it lands Heads for the first time, and the payoffs grow in magnitude without bound. But unlike the St. Petersburg game, the Pasadena game alternates rewards with punishments according to whether n , the number of tosses required, is odd or even; and its payoffs grow in absolute value as $2^n/n$ utiles. So with probability $\frac{1}{2^n}$ the Pasadena game pays $\$(-1)^{n-1}2^n/n$.

As a result, its expectation is

$$1 - \frac{1}{2} + \frac{1}{3} - \frac{1}{4} + \dots$$

respecting Open-mindedness, they are hardly reasons to question the norm; rather, they are reasons to question the theory. After all, they are *difficulties*—*unwelcome* consequences of the theory. The norm is still compelling.

This series conditionally converges (converges, but not absolutely).⁵ Consequently, expected utility theory judges the desirability of the game to be *undefined*. There is a good reason for this verdict. After all, the sum of the series is sensitive to the order of its terms. Indeed, we know from the Riemann rearrangement theorem that by suitably reordering the terms we can make the resulting series converge to any real number that we like; or we can make it diverge to ∞ or to $-\infty$; or we make it simply diverge. Yet there is no privileged ordering of the terms. (A given ordering of the terms corresponds to a given ordering of the columns of the decision matrix—but clearly the decision problem remains invariant under permutation of the columns.) So to the question ‘how good is the game’, the theory gives you no answer—just a big question mark.

But there is worse. Take some perfectly ordinary decision that you face—let it be taking a dollar, or taking two dollars. Offhand, this choice couldn’t be easier. However, if you survey your possible futures carefully, you will see that ‘taking the dollar’ future branches into two possibilities: you take the dollar and then play the Pasadena game, and you take the dollar and then do not play the Pasadena game. To be sure, the former possibility is extraordinarily improbable, and you should treat it with an appropriately scornful assignment of credence. But how much scorn does it deserve? You may be tempted to zero it out altogether—but the norm of open-mindedness will not let you. After all, your total evidence is compatible with your playing the Pasadena game. And so open-mindedness requires a positive probability from you—it may be extraordinarily small, to be sure, but positive nonetheless. However, now the damage is done: you are regarding taking the dollar as a mixture of

⁵ See Nover and Hájek (2004) and Hájek and Nover (2006) for fuller presentations of the payoff and probability tables of both the Pasadena and St. Petersburg games, and further discussion of their expectations.

two possible futures, one of which has undefined expectation, a big question mark. And a weighted average of a big question mark and anything else is a big question mark. The upshot is that you do not value the option of taking the dollar after all! The same goes for the option of taking the two dollars.⁶

You may be tempted to appeal to dominance reasoning to resolve the issue: whatever the subsequent future may bring, \$2 is better than \$1, so you are better off choosing the former. Maybe so, although notice that you in doing so you are *not* maximizing expected utility (one question mark is not larger than the other). And this reply assumes that the choice between the dollar amounts is independent of whether you play the Pasadena game or not, which may or may not be the case. If it is the case, then let's simply change the example to one in which dominance reasoning doesn't apply. You can choose between pizza and Chinese for dinner. Each option's desirability depends on how you weigh probabilistically various scenarios (burnt pizza, perfectly cooked pizza ..., over-spiced Chinese, perfectly spiced Chinese ...) and the utilities you accord them. Let us stipulate that neither choice dominates the other, yet it should be utterly straightforward for you to make a choice. But it is not if the expectations of pizza and Chinese are contaminated by even a miniscule assignment of credence to the Pasadena game. If the door is opened to it just a crack,

⁶ An important recent paper by Terrence Fine (2008) prompts us to state this point a little more carefully, but not in a way that matters to the point. Fine shows that consistent with the preference axioms of standard decision theory, the Pasadena game can be valued at any real number whatsoever. Moreover, the *Altadena* game, in which all of the Pasadena game payoffs are increased by a dollar, can be valued at any real number whatsoever, *independently of the value assigned to the Pasadena game*. Then although two dollars followed by the Pasadena game is equivalent to one dollar followed by the Altadena game (their payoff and probability profiles are identical), the former can be valued differently from the latter, by any amount. So our point can be restated: decision theory places absolutely no constraint on the values assigned to the option of taking the dollar and taking two dollars once the Pasadena game gets mixed into their evaluation.

it kicks the door down and swamps all expected utility calculations. You cannot even choose between pizza and Chinese.

This problem generalizes to any decision, however mundane. The undefined expected utility of the Pasadena game swamps the expected utility of any option that you face in just the same way—replace ‘pizza’ and ‘Chinese’ by any two options. The upshot is that the norms of Open-mindedness and of Some Expected Utility Maximization apparently conflict. If you assign any positive probability whatsoever to the Pasadena game, then all your options have undefined expectation; you violate Some Expected Utility Maximization, a sin against practical rationality. But if you zero out the Pasadena game, then you violate open-mindedness, a sin against theoretical rationality. We submit that this is a paradox for our theory of rationality.

One way to banish the paradox would be to banish the game. Since the Pasadena’s vexing expectation is the result of a probability function and a utility function working in tandem, these functions are the obvious targets.

On the side of probability, a reply that we have repeatedly heard is that you really should *zero out* the Pasadena game—if you don’t, then you deserve all the trouble that you get. This means either that your total evidence does rule out the Pasadena game, or that Open-mindedness is not a norm. Now, if your total evidence rules out the game, it is a lot richer than ours! But we don’t believe that; whatever your evidence—experiences, memories, testimony of others—may be about coin tossing, monetary amounts, and the possibilities for pairing the two, we can be sure that it does not *entail* your never playing the Pasadena game. So you must insist that rationality permits assigning probability 0 to something compatible with your total evidence. Among other things, on the betting interpretation of credence this means betting at

any odds against a contingent proposition concerning which you have only limited evidence. Fingers crossed!

On the side of utility, the obvious move is to insist that all utility functions are bounded—see e.g. Hardin (1982) and Jeffrey (1983), among others. Hájek and Nover rebut this move at length in their papers; a quick rebuttal will have to suffice here. To be sure, humans plausibly have finite capacities for punishments and rewards—there's a limit to how much a person can abhor or adore. But Bayesian decision theory is up to its neck in idealization, and it is supposed to apply to rational agents in the abstract. Indeed, it assumes that the agents have infinite capacities in other respects—for example, that they are logically omniscient. Above all, it typically assumes its agents to have *credences with unbounded sharpness*—every assignment is sharp to infinitely many decimal places. It would be a double standard for the Bayesian to insist that agents should be infinitely idealized when it comes to credences, but to plead psychological realism when it comes to utility functions. Yet once we idealize the utility functions also, imposing a bound on them seems ad hoc. So we remain unpersuaded by this move, and we believe the paradox remains alive.

We have regarded Open-mindedness as a norm, but we need not do so in order to generate a paradox. Suppose that you do not *want* to assign (sharp) probability 0 to any possibility that your total evidence does not rule out, for whatever reason—not necessarily because you regard Open-mindedness as a norm, but just because you happen to *like* being open-minded (in this doxastic sense). It seems that this is at least a rationally *permissible* way for you to be, even if it is not rationally required. Still, the paradox will have you in its grip. For you will still apparently come into conflict with the rational *requirement* of Some Expected Utility Maximization.

Indeterminate credences to the rescue! A way of being both open-minded and maximizing expected utility is to have indeterminate credence for the Pasadena game: your lower probability is 0, and your upper probability is some number greater than 0. (It does not matter which, provided it is no greater than 1 of course; it can be as small as you like, provided it is positive.) The lower probability ensures that you are able to make at least some decisions by maximizing expected utility; the upper probability ensures that you are open-minded, not assigning (sharp) probability 0 to the Pasadena game. We may thus distinguish you from all orthodox Bayesian agents, who either are unable to use expectations to make any decision, or who treat the Pasadena game as doxastically on a par with a logical contradiction, when clearly it is not.

There are various decision rules that we might employ that cater to indeterminate credences. Let us explain one that we find appealing (but once again, our main concern is with the rationality of indeterminate credences rather than the details of any particular formulation of them, and that includes their associated decision theory). If your credence is indeterminate over the interval $[a, b]$ for proposition X , then for each number r in $[a, b]$, there is a sharpening of your credence that assigns X probability r . More generally, we might allow your credence to be indeterminate over other sets. We may allow your utilities to be indeterminate also. Your overall mental state, then, is represented by a set of \langle utility function, credence function \rangle pairs, which we will call your *representational set*.

So much for your mental state; what about your decisions? We may formulate a generalized decision rule of expected utility maximization as follows:⁷

For each $\langle V, C \rangle$ in your representational set: you are rationally required to Φ according to $\langle V, C \rangle$ iff Φ -ing maximizes expected utility according to $\langle V, C \rangle$.
 You are rationally required to Φ (outright) if
 1) for some $\langle V, C \rangle$ you are rationally required according to Φ according to $\langle V, C \rangle$

⁷ Thanks here to Jonathan Schaffer for help with streamlining this formulation.

and

2) for no $\langle V', C' \rangle$ are you rationally required to do something other than Φ according to $\langle V', C' \rangle$.

Here's the intuitive idea. For each option, calculate its expected utility according to a given \langle utility function, credence function \rangle pair, $\langle V, C \rangle$, and note the winning option if there is one (as voted for by this pair). Do this again for another pair, $\langle V', C' \rangle$. And so on. Some pairs may yield an undefined expectation, which we may regard as an 'informal' vote (as we say in Australia)—since no preference has been lodged, the vote is discarded. Once all the genuine votes are in, check to see if there is a unanimous winner. If there is, this is the action that you are rationally required to perform.

Let's see how an indeterminate credence for the Pasadena game—for definiteness, let it be $[0, 1/1,000,000]$ —allows you to maximize expected utility. It simplifies matters to suppose that your utilities are sharp, although this is not essential. What should you choose: pizza or Chinese? Sharpening with a probability of 0 for the Pasadena game, it no longer contaminates your expectation calculations—it's as if the Pasadena game were never there in the back of your mind. So your problem reduces to an ordinary comparison of the expected utilities of pizza and Chinese. Let's stipulate that pizza wins. Sharpening with any positive probability from $(0, 1/1,000,000]$ for the Pasadena game yields garbage in the expectation calculations—its undefined expectation swamps the calculations. All you get is a big question mark, and no vote either way. Now collect the votes. You either get votes for pizza, or 'informal' votes. Discard the latter. Unanimity for pizza remains, so it wins the election. You are not paralyzed after all: go eat some pizza!

Nor are you closed-minded. For the positive part of your indeterminate credence ensures that the Pasadena game *is* there in the back of your mind. We can easily tell

you apart from an agent whose credence function dogmatically zeroes out the Pasadena game: you will not sell a million-dollar bet on the Pasadena game's occurrence for less than a dollar, while that agent will give it away for free!

We have argued that rationality requires an assignment to the Pasadena game that is indeterminate over a set that includes 0. But as always, there are ...

Objections

For the sake of brevity, we will consider just two natural objections. We realize that there are doubtless other objections that merit our attention. But if problems for our solution remain, we predict that they will be at least as serious problems for orthodox Bayesianism too. For we see no painless solution to our dilemma, and in particular no such solution that maintains sharp probabilities.

*First objection.*⁸ “Your solution does not uphold the spirit of open-mindedness. It is only thanks to the *closed*-mindedness of the lower probability of 0 to the Pasadena game that your agent overcomes paralysis (in deciding between pizza and Chinese, etc.). True open-mindedness would banish the 0 altogether.”

We reply: We may clearly distinguish our agent from one who assigns a sharp 0 to the Pasadena game. Firstly, there's the simple point that an indeterminate assignment is importantly different from a sharp one. The sharp agent's assignment to the Pasadena game is indistinguishable from her assignment to a contradiction; this is not true of our agent. Secondly, the agents will have different updating dispositions. For example, if both update by conditioning on their evidence (our agent doing so, say, but conditioning each probability function in her representational set), then the sharp

⁸ We are grateful to Aidan Lyon and Susanna Rinard for suggesting versions of this objection.

agent's assignment remains unchanged whatever the evidence,⁹ while our agent's may change. Thirdly, while both agents may behave the same way with respect to the pizza versus Chinese decision, their behaviours on other decisions will betray their different mental states. For example, the sharp agent will bet against the Pasadena game at any odds, while ours will not.

Second objection: You might gain evidence, on the basis of which your lower probability for Pasadena game is no longer 0. For example, you might learn that the result of a fair coin toss determines whether or not the Pasadena game will be played. You might even learn that you are in fact playing the game. Or you might gain evidence that, while keeping your credence indeterminate, should lift your lower probability above 0. In that case, your state of obedience to the norms of open-mindedness and expected utility maximization is precarious, hostage to your not getting such evidence. Should you pay money to avoid getting it?!

We reply: We take this objection very seriously, but think it can be defused, or at least mitigated, nonetheless. Firstly, it is not completely obvious that you could get evidence of the kind imagined. To take the hardest case first: how could you learn that you are in fact playing the Pasadena game? You see a coin being tossed repeatedly; you see money changing hands; ... All of this is compatible with your playing some finite truncation of the Pasadena game, or another infinite game with the same initial profile of pay-offs, rather than the genuine game. So whatever you see, you should retain some scepticism that you are really playing the game, and this scepticism is

⁹ Here we assume that *orthodox* Bayesianism is the alternative to our proposal. To be sure, an *unorthodox* Bayesianism that appealed to primitive conditional probability functions, such as Popper functions, could allow an assignment of 0 to be raised by a suitably redefined notion of conditionalization. We continue to assume orthodox Bayesianism as our foil in our second reply to the second objection, coming shortly.

well represented with the 0 in your interval of indeterminacy. All the more you should retain some scepticism, and retain the 0 in your interval of indeterminacy, if you are told that a coin toss will determine whether or not you will play the game, or if you get evidence that keeps your credence indeterminate.

This brings us to our second reply: those who insist that one should *zero out* the Pasadena game altogether can hardly take the moral high ground here. Indeed, their position is strictly worse off than ours with regard to this objection: while we at least give the agent the option of *raising the upper probability* that she assigns to the Pasadena game in response to evidence, the zero-outer is permanently trapped at zero. Moreover, she clearly violates the norm of open-mindedness, so she is not offering a solution to the dilemma.

Thirdly, an interpretivist might insist that a rational agent's credal state is indeterminate in the way that we have described, because no other attribution from an ideal interpreter would *rationalize* her dispositions.¹⁰ Attributing to her a sharp zero for the Pasadena game would not do so: it would portray her as closed-minded. Nor would attributing to her a sharp positive probability do so: it would portray her as paralyzed by the lights of decision theory. The rationality *norms* of Open-mindedness and Some Expected Utility Maximization become *constraints* on any rationalizing interpretation of her. Thus, according to the interpretivist, and assuming her rationality, they become constraints on what her credal state could be.

But what if an unimpeachable source—God, if you like—tells you that this coin toss will determine whether or not you will play the Pasadena Game? Then this should be reflected in your behavior—or lack thereof. If you believe the source, by the lights of standard decision theory you should suddenly be unable to decide

¹⁰ Thanks here to Adam Elga.

between any pair of options, and you should remain undecided even if one of these options is arbitrarily sweetened. If, however, the interpreter sees you happily ordering pizza, then she must find a way to attribute to you some zeroing out of the Pasadena game in your credal state. (If she attributes credence $\frac{1}{2}$ to you for the Pasadena game, your behavior makes no sense.) In that case it seems that you are not treating the source as unimpeachable after all—your decisive behaviour betrays you.

Finally, the phenomenon of it apparently being rational to pay to avoid cost-free evidence, while somewhat puzzling, is a familiar one—see Kadane et al. (1996) on how (merely) finitely additive probabilities can give rise to this phenomenon, owing to their non-conglomerability. Closer to home, Seidenfeld and Wasserman (1993) suggest that the phenomenon may be found in cases of *dilation*, in which the interval of indeterminacy for some proposition is strictly contained in the corresponding interval after conditionalization on *any* member of some partition. To be sure, there seems to be something curious about such cases (although Seidenfeld and Wasserman show how widespread the phenomenon of dilation can be). But then, we described the Pasadena game with exactly that word ourselves.

Conclusion

We claim to have found a way for you to serve both your masters of rationality simultaneously: you are able to rationalize your choices in ordinary decision problems (such as pizza versus Chinese), while staying open-minded about the Pasadena game. No sharp probability assignment can allow you to do this. We offer this as an argument that indeterminate credences play a role for the rational agent that sharp probabilities cannot.

More generally, we have we have offered three new arguments for taking seriously indeterminate credences. We think that the argument from interpretivism shows that rationality *permits* indeterminate credences, although that it is reason enough to take them seriously. We think that the argument from indeterminate chances shows that rationality permits, and perhaps even *requires* indeterminate probabilities, and this is a reason to take them *very* seriously. We do not claim that any of these arguments clinches the case for indeterminate credences; as usual, there is plenty of room for debate. Maybe each argument that we have considered could be answered in a different way by friends of sharp credences. (For example, Easwaran's (2008) very different way of handling the Pasadena game allows one to maintain sharp credences.) Rather, our strategy has been to try to build the case for indeterminate credences from a number of angles. Indeterminate credences provide a unified way of dealing with the various issues and problems that we have raised.

While the study of *belief* is a staple in the philosophy of mind and epistemology, *degrees of belief* are an important and under-explored topic. (See Eriksson and Hájek 2007 for more on this theme.) And indeterminacy is an important and under-explored topic in the study of degrees of belief. We thus welcome philosophers of mind and epistemologists to join us and other friends of indeterminate degrees of belief in their study.

REFERENCES

- Bernstein, A. R., and F. Wattenberg (1969): "Non-standard measure theory", in W. A. J. Luxemburg, ed., *Applications of Model Theory of Algebra, Analysis, and Probability*. New York: Holt, Rinehard and Winston.
- Easwaran, Kenny (2008): "Strong and Weak Expectations", *Mind* 117(467):633-641

- Edwards, W., Lindman, H., and Savage, L. J. (1963): “Bayesian Statistical Inference for Psychological Research”, *Psychological Review* 70, 193-242.
- Ellsberg, Daniel (1961): “Risk, Ambiguity and the Savage Axioms”, *Quarterly Journal of Economics* 75, 643–669.
- Eriksson, Lina and Hájek, Alan (2007): “What Are Degrees of Belief?”, *Studia Logica* 86, No. 2. (July), 183-213.
- Fine, Terrence L. (2008): “Evaluating the Pasadena, Altadena, and St Petersburg Gambles”, *Mind* 117 (467), 613-632.
- Hájek, Alan (2003): “What Conditional Probability Could Not Be”, *Synthese* 137, No. 3 (December), 273-323.
- Hájek, Alan and Harris Nover (2006): “Perplexing Expectations”, *Mind* 115 (July), 703-720.
- Hájek, Alan and Harris Nover (2008): “Complex Expectations”, *Mind* 117 (July), 643-664.
- Hardin, Russell (1982): *Collective Action*, Baltimore: The Johns Hopkins University Press.
- Jeffrey, Richard (1983): *The Logic of Decision* (2nd ed.), Chicago: University of Chicago Press.
- Jeffrey, Richard (1992): *Probability and the Art of Judgment*, Cambridge: Cambridge University Press.
- Kadane, Joseph B, Mark J. Schervish, and Teddy Seidenfeld (1996): “Reasoning to a Foregone Conclusion”, *Journal of the American Statistical Association* 91, No. 435 (September), 1228-1235.
- Levi, Isaac (1982): “Conflict and Social Agency”, *Journal of Philosophy* 79, 231-241.

- Levi, Isaac (2000): "Imprecise and Indeterminate Probabilities", *Risk, Decision and Policy* 5, 111-122.
- Lewis, David (1974): "Radical Interpretation", *Synthese* 23, 331-44.
- Lewis, David (1980): "A Subjectivist's Guide to Objective Chance", in *Studies in Inductive Logic and Probability*, Vol II., ed. Richard C. Jeffrey, University of California Press.
- Lewis, David (1994): "Humean Supervenience Debugged", *Mind* 103, 473-490.
- Nover, Harris and Alan Hájek (2004): "Vexing Expectations" *Mind* 113 (April), 237-249.
- Papamarcou, Adrianos and Terrence L. Fine (1986): "A Note on Undominated Lower Probabilities", *The Annals of Probability* 14, No. 2, 710-723.
- Seidenfeld, Teddy, Mark J. Schervish, and Joseph B. Kadane (1989): "On the Shared Preferences of Two Bayesian Decision Makers", *The Journal of Philosophy* 86, 225-244.
- Seidenfeld, Teddy and Larry Wasserman (1993): "Dilation for Sets of Probabilities", *The Annals of Statistics* 21, No. 3 (September), 1139-54.
- Smithson, Michael (1999): "Conflict Aversion: Preference for Ambiguity vs. Conflict in Sources and Evidence", *Organizational Behavior and Human Decision Processes* 79, 179-198.
- Sorensen, Roy (2001): *Vagueness and Contradiction*, Oxford: Oxford University Press.
- Walley, Peter (1991): *Statistical Reasoning with Imprecise Probabilities*. London: Chapman Hall.

Walley, Peter (1996): “Inferences from Multinomial Data: Learning about a Bag of Marbles” (with discussion), *Journal of the Royal Statistical Society, Series B* 58, 3-57.

Walley, Peter and Terrence L. Fine (1982): “Towards a Frequentist Theory of Upper and Lower Probability”, *Annals of Statistics* 10, 741-761.

Williamson, Timothy (1994): *Vagueness*, London: Routledge.