Probability in Philosophy

Brian Weatherson

Rutgers and Arché

June, 2008

Credences are Probabilities

Orthodox Bayesianism says that probability theory matters because (rational/coherent) credences are probabilities

- We get different Bayesian theories depending on how we interpret rationality or coherence

What Credences Are

Credences are basically degrees of confidence

- The more confident you are in p, the higher your credence in p
- Credences are sometimes described as degrees of belief, but this is perhaps misleading
- But there can be two propositions, p and q, such that I think both are quite unlikely, but my credence in p is higher than in q

What Probabilities Are

- Among Bayesians, the big dispute is over countable additivity
- Also there is some dispute over whether credences are fundamentally conditional or unconditional
- But we'll stick to cases involving finitely many propositions, so neither of these complications are relevant

Update is by Conditionalisation

As well as constraints on what credences at a time are, orthodox Bayesianism has constraints on how to update credences

- The short version is we update by *conditionalisation*
- That is, when you get evidence E, the new credence of H is the old credence of H given E
- This holds for any H

Traditional Conceptions of Evidence

What is it to get evidence E?

- The standard Bayesian picture is neutral on this question
- But most theorists who developed Bayesianism had a very internalist picture of evidence
- So evidence was thought of as being something like sense-data
- But note that E has to be a proposition for it to be something you conditionalise on

More Sophisticated Conceptions of Evidence

It's no part of the Bayesian theory of evidence that evidence is internal

- But as far as I know very little work has been done on working out the details
- There is, I think, a potentially enormous research project here
- Apart from Williamson, there is very little work done on probabilistic epistemology by people who are not, by inclination, more internalist than the majority of epistemologists

Evidence and Certainty

Note it's a theorem that $Pr(E|E) = 1$

- So updating on E takes the probability of E to 1
- If updating by conditionalisation is the only way to update, it follows that you can only get grounds to change credences if you become certain in something
- Arguably, that's not epistemologically plausible

Evidence as Indefeasible

Note also that if $Pr(E) = 1$ and $Pr(E') > 0$, then $Pr(E|E') = 1$

- So if you conditionalise on E, and then on E', you're still certain that E
- So if E is evidence, it can't be defeated
- Or if it can, it's only by a probability 0 event
- Again, this isn't particularly plausible

Jeffrey Conditionalisation

Jeffrey Conditionalisation is a response to this. The idea is that we force the probability of E to be some value x, and the new credence function is given by the following formula.

$$Pr_{new}(A) = x \times Pr_{old}(A|E) + (1 - x) \times Pr_{old}(A|\neg E)$$

Standard conditionalisation is the special case where $x = 1$.

Jeffrey Conditionalisation

One objection to Jeffrey Conditionalisation (more precisely, to the idea that it is epistemologically useful) is that it is unclear where the x comes from.

- If I look down College Street, and see someone who looks a bit like Crispin walking south, I should increase my credence that Crispin is on College Street
- But how can I tell whether I should increase it to 0.4, to 0.6, to 0.8 etc?
- There is a threat of a 'false precision' problem here

Jeffrey Conditionalisation

Many people also worry that Jeffrey Conditionalisation is not symmetric

- Making $Pr(E_1) = x$ then making $Pr(E_2) = y$ won't in general have the same output as making $Pr(E_2) = y$ then making $Pr(E_1) = x$
- To see this formally, let $E_1 = E_2$ and $x \neq y$
- Perhaps this is a feature, not a bug, if we want evidence to be sensitive to background
- Although it might cause problems if we simultaneously get two distinct pieces of (partial) evidence

Jeffrey Conditionalisation

We can put these two worries together

- Classic conditionalisation, for better or worse, gave us a somewhat clear separation of what was learned from what was background
- What you do with what you learn is sensitive to background, but what you learn is not
- Jeffrey Conditionalisation doesn't have that; what my new credence that Crispin is on College Street should be, if I get a glimpse of someone who looks like Crispin, depends a lot on my background
- That's why (a) it's hard to say what value this x should take, and (b) Jeffrey conditionalisation is not symmetric

Coherence and Rationality

Bayesians take it as an important result that credences are coherent only if (perhaps if and only if) they form a probability function. What does coherence mean here?

- It doesn't mean *rational*
- Someone whose credences are only defined over $p, \neg p, p \vee \neg p, p \wedge \neg p$, and has $Pr(p) = Pr(p \vee \neg p) = 1$ and $Pr(\neg p) = Pr(p \wedge \neg p) = 0$ is coherent, more so than you or I
- But if $p =$ *The moon is made of green cheese*, they are less rational

Two Relations Between Belief and Logic

Here are two logical properties your belief set might have.

Consistency	Your beliefs are logically consistent
Closure	Your beliefs are closed under entailment

Having credences be probabilities is meant to be a kind of logical coherence. But it isn't clear, as we'll see, whether it is more akin to **Consistency** or **Closure**.

Assumptions about Betting Behaviour

- Say a $p-$bet is a bet that returns \$1 if p is true, and 0 otherwise.
- Assume that each dollar is worth as much to you as the previous dollar.
- Let Cr be your personal credence function.

Then plausibly

- The value to you of a $p-$bet is $\$Cr(p)$.
- If offered a $p-$bet for less than $\$Cr(p)$, you'll buy it, and you'll sell a $p-$bet for anything above that.

Violations Lead to Sure Loss

Assume Cr is not a probability function.

- This can happen in two ways
- First, your credence in some logical truth A might be less than 1
- Second, you might violate the addition postulate, so $Cr(A) + Cr(B) \neq Cr(A \vee B)$, with disjoint A, B

Either way, there will be a series of bets you'll take that will lead to sure loss

Violations Lead to Sure Loss

Assume $Cr(A) = x < 1$, where A is a logical truth

- Consider someone who offers to buy from you a bet on A for $\$\frac{x+1}{2}$
- This is greater than x, so you should sell
- But the bet is sure to win, so you'll have to pay \$1, and you received less than \$1

So you'll make a sure loss

Violations Lead to Sure Loss

Assume $Cr(A) + Cr(B) = Cr(A \vee B) - x$, where A and B are disjoint, and $x > 0$.

- Consider someone who offers to buy from you a bet on A for $\$Cr(A) + \frac{x}{4}$, and on B for $\$Cr(B) + \frac{x}{4}$
- They then offer to sell you a bet on $A \vee B$ for $\$Cr(A \vee B) - \frac{x}{4}$
- All these trades are good value to you, so you should make them
- But you'll have sold just as many winning bets as you bought, i.e. one bet if $A \vee B$ is true and none otherwise
- And your cash balance will be down $\$\frac{x}{4}$

So you'll make a sure loss

Violations Lead to Sure Loss

Assume $Cr(A) + Cr(B) = Cr(A \vee B) + x$, where A and B are disjoint, and $x > 0$.

- Consider someone who offers to sell you a bet on A for $\$Cr(A) + \frac{x}{4}$, and on B for $\$Cr(B) + \frac{x}{4}$
- They then offer to buy from you a bet on $A \vee B$ for $\$Cr(A \vee B) - \frac{x}{4}$
- All these trades are good value to you, so you should make them
- But you'll have sold just as many winning bets as you bought, i.e. one bet if $A \vee B$ is true and none otherwise
- And your cash balance will be down $\$\frac{x}{4}$

So you'll make a sure loss

No Sure Losses if You Comply

- We won't try to prove this here, but it can be shown that if *Cr* is a probability function, then there is no series of bets that leads to sure loss in this way
- So you are vulnerable to sure loss iff *Cr* is not a probability function

Why These Losses are Bad

Of course, anyone could lose a bet. We need more than that to say that you are *incoherent*

- The idea is that if you are subject to sure loss, and you could in principle know this, then you should correct your credences
- A coherent agent would not subject herself to *sure* loss in this way

Why These Losses are Bad

Of course, anyone could lose a bet. We need more than that to say that you are *incoherent*

- The idea is that if you are subject to sure loss, and you could in principle know this, then you should correct your credences
- A coherent agent would not subject herself to *sure* loss in this way

That's how the argument is usually put, though presumably knowable loss is more philosophically important than sure loss. Having credence < 1 in *Water is* H_2O leads to sure loss, but isn't incoherent.

Dynamic DBAs

In general, a *Dutch Book Argument* is an argument that a certain credal state leads to a sure loss, and hence is incoherent. We've looked at the most important *static* Dutch Book Argument, but there are also two important *dynamic* DBAs.

1. The Lewis/Teller argument for conditionalisation
2. The van Fraassen argument for Reflection

DBA for Conditionalisation

Assume $Cr(H|E) = x$, but you plan to have $Cr(H) = y < x$ if you learn E. (The case where $y > x$ is symmetric, and we'll ignore it.)

- Assume also, and this is a little controversial, that the agent will know whether or not her evidence is E
- It's worth thinking about how this effects the argument
- Also say $Cr(E) = z$

DBA for Conditionalisation

Assume $Cr(H|E) = x$, but you plan to have $Cr(H) = y < x$ if you learn E. (The case where $y > x$ is symmetric, and we'll ignore it.)

- Assume also, and this is a little controversial, that the agent will know whether or not her evidence is E
- It's worth thinking about how this effects the argument
- Also say $Cr(E) = z$

Note that in what follows a *conditional bet* on p conditional on q will be a bet that pays \$1 if $p \wedge q$ is true, loses if $p \wedge \neg q$ is true, and for which the purchase price is returned if $\neg p$ is true.

DBA for Conditionalisation

The bookmaker starts off by selling you a bet on H conditional on E for \$$x$, and the following bet

- If E is true, you get $\$(1 - z)(x - y)$
- If E is false, she gets $\$z(x - y)$
- In effect, it's a bet on E at with the stakes being $\$(x - y)$ rather than \$1

DBA for Conditionalisation

The bookmaker starts off by selling you a bet on H conditional on E for $\$x$, and the following bet

- If E is true, you get $\$(1 - z)(x - y)$
- If E is false, she gets $\$z(x - y)$
- In effect, it's a bet on E at with the stakes being $\$(x - y)$ rather than $1

If E is false she quits while ahead $\$z(x - y)$. If E is true, she then sells you a bet on $\neg H$ for $\$1 - y$.

Now you've paid $\$1 + x - y$ for bets on H and $\neg H$, exactly one of which will win, so you'll end up down $\$x - y$. You also will win the bet on E, but that will return $\$(1 - z)(x - y)$, so you end up down $\$z(x - y)$.

DBA for Reflection

Reflection is the following principle, where Cr_0 and Cr_1 are your credences at an earlier and later time

Reflection $Cr_0(H|Cr_1(H) = x) = x$

It is very implausible, especially if you have good evidence at the earlier time that you'll be irrational at the later time. But it can be supported by the same kind of argument. That is, if $Cr_0(H|Cr_1(H) = x) = y < x$ there is a Dutch Book that can be made against you. (Or, for that matter, if $y > x$.

DBA for Reflection

I won't, in the interests of time, go through the details, but the structure is the same

- At the first time the bookmaker offers you a bet on H conditional on $Cr_1(H) = x$, and a low stakes bet that $Cr_1(H) \neq x$
- At the later time, if $Cr_1(H) = x$, the bookmaker offers you a bet on $\neg H$ at worse odds, or, equivalently, offers to buy the bet you have on H back at worse odds

Until very recently, it seemed these arguments were as good as each other. This has been challenged in some forthcoming papers, and perhaps this consensus will change in future years.

Agents Who Don't Know Their Own Credences

The Dutch Book Argument was meant to reveal that there was a sure loss I knew I was vulnerable to. But imagine the following three things are true of me.

1. My credence in p is 0.7
2. My credence in $p \vee q$ is 0.6
3. My 'margin of error' on my credences is 0.2. That is, if my credence in a proposition is x, I at most know it is in $[x - 0.2, x + 0.2]$

Then there's no way to make a Dutch Book against me without exploiting knowledge (of my credences) that I don't have, and couldn't have. And that someone with more knowledge can sell me bets that will lose is no sign of *incoherence* on my part.

Use Value and Exchange Value

The assumption about betting behaviour isn't, as far as I can tell, followed by the people who are my paradigm at least of successful bettors, namely professional bookmakers

- Professional bookmakers will make bets even if they don't think they are winning bets, in some sense, as long as they think they can sell them at a profit
- This is an empirical claim that could use more defence, but I'm pretty sure it's true for various reasons

I conclude that the *use value* of a bet on p, to use an old Marxist term, is $Cr(p)$, but whether I should buy or sell the bet might depend also on its *exchange value*, and that could come apart from its use value.

Use Value and Exchange Value

In principle, as long as there is any more trading to do, or at least I have a non-zero credence that there is more trading to do, I might buy or sell bets depending on their exchange value

- But the Dutch Book argument assumes that (a) there is more betting coming, but (b) I'll trade bets for their use value
- If I don't know that there is more betting coming, then the bookmaker is exploiting knowledge I don't have, and hence my loss is no sign of incoherence
- If I do know there is more betting coming, then arguably the losses I end up incurring are a sign of my poor trading strategies, not my allegedly incoherent credences.

Use Value and Exchange Value

In principle, as long as there is any more trading to do, or at least I have a non-zero credence that there is more trading to do, I might buy or sell bets depending on their exchange value

- But the Dutch Book argument assumes that (a) there is more betting coming, but (b) I'll trade bets for their use value
- If I don't know that there is more betting coming, then the bookmaker is exploiting knowledge I don't have, and hence my loss is no sign of incoherence
- If I do know there is more betting coming, then arguably the losses I end up incurring are a sign of my poor trading strategies, not my allegedly incoherent credences.

I don't know whether this is the same objection as the one (due largely to Isaac Levi) that the smart agent will 'see the Dutch Book coming'.

Wrong Kind of Flaw

The Dutch Book Argument was meant to reveal a kind of doxastic incoherence

- But what we've been left with here is at worst a *practical* flaw
- Perhaps there's a further argument that this practical flaw is distinctively revealing of a doxastic incoherence, but that needs to be shown, and hasn't been.

No Actual Chance of Loss

And it isn't really that much of a practical flaw.

- There aren't that many Dutch Bookies around!
- And if there are, you can avoid them by all sorts of means
- For instance, you can simply not reveal what your credences are

Undecidable Propositions

It really isn't a practical flaw if one of the propositions involved is undecidable.

- The argument assumes that every bet will be settled
- But some of the strongest objections to Bayesian probabilism concern undecidable propositions
- It's really unclear what kind of flaw the Dutch Book Argument *could* reveal about those propositions

Selling Bets You Don't Have

The next two objections turn on the assumption about betting behaviour

- What is it to *sell* a bet on A for \$$x$, when you don't have such a bet
- You might think it just is to *buy* a bet on $\neg A$ for \$$(1 - x)$
- But if we think of the trades that way, the argument breaks down
- To see this, consider a case where $Cr(A) = 0.8, Cr(\neg A) = 0, Cr(A \vee \neg A) = 1$
- The algorithm for constructing a Dutch Book asks you to sell a bet on A for 85c. But you wouldn't buy a bet on $\neg A$ for 15c.

Complementary Goods

One very important concept from economics is that some goods are *complements* and some goods are *substitutes*.

Complements: Goods a and b are complements if b is more valuable when you also have a

Substitutes: Goods a and b are substitutes if b is less valuable when you also have a

- So a CD player and a CD are complements
- While two copies of the same CD are substitutes

Complementary Goods

Now it's an interesting question which bets are complementary or substitute goods

- It's agreed on all sides that if the marginal value of money is not stable, then many bets are complements
- For instance, if you really can't afford to lose money, you might not be prepared to spend more than $300 for a bet that pays $1000 if this fair coin comes up heads
- And you might not pay more than $300 for the same bet on it coming up tails
- But you'd pay much more than $600 for the pair of bets

What's harder to say is whether the changing marginal value of money is the only cause of complementarity

Complementary Goods

In practice, most of the non-probabilist approaches to credence (which we'll look at next week) are committed to various bets being complements

- So imagine $Cr(p \vee q) > Cr(p) + Cr(q)$, with p and q disjoint
- Then if you own a bet on p, the value of a bet on q will not be $Cr(q)$
- It rather will be $Cr(p \vee q) - Cr(p)$
- And given that value, you won't be subject to any Dutch Book

I tend to think this is a strong objection to DBAs, though I think this is somewhat of a minority view

Summary

We looked at seven objections to DBAs

1. Agents Who Don't Know Their Own Credences
2. Use Value and Exchange Value
3. Wrong Kind of Flaw
4. No Actual Chance of Loss
5. Undecidable Propositions
6. Selling Bets You Don't Have
7. Complementary Goods

These look like substantial objections!

Christensen's Idea

In response to some of these objections, David Christensen has been pushing a 'depragmatised' DBA

- The idea is that we shouldn't take the DBA to be a sign that someone with non-probabilist credences will really lose money
- That might not be true, and even if it were, it wouldn't necessarily show their credences were incoherent

Rather, the idea is that the DBA reveals that a non-probabilist has incoherent valuations of bets

Christensen's Idea

Let's work through this with an example where p and q are disjoint.

- $Cr(p) = 0.2$
- $Cr(q) = 0.3$
- $Cr(p \vee q) = 0.7$

Then the agent values bets in the following way.

Bet on p	20 cents
Bet on q	30 cents
Bet on $p \vee q$	70 cents

Christensen's Idea

But a bet on $p \vee q$ just is a bet on p plus a bet on q, so the agent should value it the same way.

- The incoherence is that she values a bet on $p \vee q$ at 70 cents presented one way (i.e. as such) and at 50 cents presented another way (i.e. a bet on p plus a bet on q)

Responding to Objections

Remember that e looked at seven objections to DBAs, all of which looked problematic on the pragmatic understanding of the argument

1. Agents Who Don't Know Their Own Credences
2. Use Value and Exchange Value
3. Wrong Kind of Flaw
4. No Actual Chance of Loss
5. Undecidable Propositions
6. Selling Bets You Don't Have
7. Complementary Goods

Responding to Objections

Christensen's way of setting out the argument seems like a good response to the first four

1. ~~Agents Who Don't Know Their Own Credences~~
2. ~~Use Value and Exchange Value~~
3. ~~Wrong Kind of Flaw~~
4. ~~No Actual Chance of Loss~~
5. Undecidable Propositions
6. Selling Bets You Don't Have
7. Complementary Goods

Responding to Objections

Christensen's way of setting out the argument seems like a good response to the first four

1. ~~Agents Who Don't Know Their Own Credences~~
2. ~~Use Value and Exchange Value~~
3. ~~Wrong Kind of Flaw~~
4. ~~No Actual Chance of Loss~~
5. Undecidable Propositions
6. Selling Bets You Don't Have
7. Complementary Goods

If anyone is worried about these concessions, perhaps we should spend more time on this point

Responding to Objections

Whether Christensen can respond to 5 depends on tricky questions about truth and decidability that we aren't addressing here

1. ~~Agents Who Don't Know Their Own Credences~~
2. ~~Use Value and Exchange Value~~
3. ~~Wrong Kind of Flaw~~
4. ~~No Actual Chance of Loss~~
5. ??Undecidable Propositions
6. Selling Bets You Don't Have
7. Complementary Goods

Responding to Objections

It seems to me though that there really isn't any response here to 6 or 7

1. ~~Agents Who Don't Know Their Own Credences~~
2. ~~Use Value and Exchange Value~~
3. ~~Wrong Kind of Flaw~~
4. ~~No Actual Chance of Loss~~
5. ??Undecidable Propositions
6. Selling Bets You Don't Have
7. Complementary Goods

Complementary Goods

There isn't anything wrong with, in your current setting, valuing each of a CD and a CD player in a way that their value is less than the value of the package

- That's just what it is for them to be complementary goods
- The issue is whether bets could be complementary goods
- Christensen, I think, wants to say that if we presuppose a constant marginal utility of money, that's ruled out
- But I don't see why we should believe this, and I've never seen a compelling argument against this possibility

Dynamic Dutch Book Arguments

It's also unclear whether it is possible to depragmatise the dynamic arguments

- In fact Christensen himself is sceptical about how good these arguments are
- So even if the objection I was making fails, it looks like there isn't a good argument here for conditionalisation

Status of Conditionalisation Rule

This does leave conditionalisation with an odd status

- On the one hand, I know of literally no direct argument for conditionalisation that's even plausibly sound
- On the other, everyone assumes in Bayesian epistemology that this is how you update
- And Bayesian epistemology can explain a lot of things in philosophy of science
- So maybe there's an indirect argument for it

Status of Conditionalisation Rule

Bas van Fraassen noted an odd point about the argument for conditionalisation

- The most it could possibly show is that you shouldn't have a plan to not conditionalise
- But that doesn't show that you have to have a plan to conditionalise
- You might have no plan at all
- Again, perhaps the indirect argument can handle this

Comparative Constraints

We'll move quickly through the next two arguments, starting with an argument from comparative constraints on preference

- Before we start thinking about numerical utility or credence, we can put a lot of coherence constraints on $\succeq$
- For instance, it should be transitive
- One argument for probabilism comes from these constraints on $\succeq$

Representability

A representation theorem for $\succeq$ says that if $\succeq$ satisfies certain constraints, then it is as if the agent is maximising expected utility relative to a probability function *Pr* and a utility function *u*.

- That is, $A \succeq B$ iff $Exp[u(A)] \geq Exp[u(B)]$, where *Exp* means 'expected value of'
- For constraints we'll be interested in, not only are there functions, but they are fairly distinctive
- *Pr* is unique
- *u* is unique up to positive affine transformation
- An affine transformation maps x onto $ax + b$, and it's positive iff $a > 0$

A Philosophical Argument

This suggests the following kind of argument

- Assume that the relevant constraints are plausible constraints on $\succeq$
- So any coherent agent satisfies them
- Then for any coherent agent, there is a function *Pr* that (a) behaves a lot like a credence function for them, and (b) is a probability function
- So (waving hands furiously here) coherent agents have probability functions for credences

Are These Really Credences

The obvious philosophical problem is the final step

- Just because there is something that behaves (in one respect) like a credence function that is also a probability function, it doesn't mean credences are probabilities
- Not even the most die-hard behaviourist should say that credences are constituted **entirely** by preferences

Purely Doxastic Agents

One way of making this vivid is to consider an agent with no preferences over worldly states

- This is a relatively extreme kind of Zen I guess
- Such an agent might have credal states, but they won't fall out of preferences

In fact, the assumptions about $\succeq$ you have to make rule this out, which raises issues about how much the constraints on $\succeq$ are really coherence constraints

Constraints on Preferences

Another batch of worries comes from looking at the details of the purported constraints on $\succeq$

- Several of the constraints are much stronger than are plausible as mere coherence constraints
- Different approaches use different constraints, but there are three major centres of concern

1. Archimedean Principle
2. Conglomerability Principles
3. Completeness Principles

Archimedes

When Ramsey originally put forward an argument of this kind, he simply said that we need an "Archimedean axiom". Different theorists do this different way, and I'll simplify things a little.

Archimedes If $A \succ B \succ C$ then there is some real number $x \in (0, 1)$ such that the agent is indifferent between B, on the one hand, and a bet that returns A with probability x, and C with probability $1 - x$

This seems unlikely if, for instance, A is you getting \$1, B is the status quo, and C is the destruction of the universe.

Lexical Utility

An axiom like this is needed to rule out *lexical* utility functions.

- These are functions that order actions on two (or more) dimensions, with the second (or further) dimensions used only as a tiebreaker
- So, for example, consider an agent who considers, for each action, the expected number of sins involved if that action is performed, and the expected level of fun if it is performed, then tries to minimise the amount of sin, and, if two actions involve the same amount of sin, perform the action that produces the most fun
- That seems coherent, but it is ruled out by this constraint

Conglomerability

We've already talked about this, so I won't say more, but these approaches require very strong conglomerability constraints

- In order to get countable additivity, we need countable conglomerability, which is problematic
- Even finite conglomerability is controversial because of the Allais and Ellsberg paradoxes
- So this is controversial

Completeness

Finally, we need a constraint to ensure that $\succeq$ is a linear order

- That is, for any A, B, we need to suppose that $A \succeq B$ or $B \succeq A$
- This isn't obviously correct

Completeness

Finally, we need a constraint to ensure that $\succeq$ is a linear order

- That is, for any A, B, we need to suppose that $A \succeq B$ or $B \succeq A$
- This isn't obviously correct

As we'll see a little later, one of the controversial things is whether Cr should satisfy $[Cr(p) \geq Cr(q)] \vee [Cr(q) \geq Cr(p)]$, and that basically has to be presupposed on this account

Scoring Rules

One interesting addition to arguments for probabilism is due to James Joyce

- He starts from considerations about how we judge credences for accuracy
- Ideally, our credence in any truth would be 1, and in any falsehood 0
- In non-ideal situations, we might try to measure how close we are to this ideal

Scoring Rules

One interesting addition to arguments for probabilism is due to James Joyce

- He starts from considerations about how we judge credences for accuracy
- Ideally, our credence in any truth would be 1, and in any falsehood 0
- In non-ideal situations, we might try to measure how close we are to this ideal

This will be hard to do in general, due to results from Tchy and others, but we can measure how close we are over a finite field of propositions.

Scoring Rules

So imagine we have some finite field $\mathcal{F}$ of propositions, closed under conjunction, disjunction and negation, such that *Cr* is defined over every $A \in \mathcal{F}$ and (this is a big assumption) is a real number.

- For each $A \in \mathcal{F}$ we look at how far $Cr(A)$ is from the ideal number, i.e. 1 for truths and 0 for falsehoods
- Let $d(A) = |Cr(A) - T(A)|$, where T is the function that maps truths to 1, and falsehoods to 0
- We have some scoring rule s, which is a function that takes A as input, and returns some 'score' as a function of $d(A)$
- So $s(A)$ might be, perhaps $d(A)^2$
- That's actually a commonly used rule

We then sum $s(A)$ over each $A \in \mathcal{F}$, and the lower that sum is, the more accurate the credences are. Call that sum the *accuracy* of *Cr*

Scoring Rules

This approach has real world applications

- For instance, something like this is how weather forecasters are judged
- Though (and I don't know if this is relevant), usually the propositions they are judged over is not closed under conjunction, disjunction or negation

Indeed, this way of judging accuracy of probabilistic predictions is a lot preferable to some older ideas based around the idea of 'calibration'

Constraints on Scoring Rules

Although the rule $s(A) = d(A)^2$ is popular, it isn't the only rule we can imagine. But we can imagine constraints on a rule.

- For instance, $s(A)$ should be a positive function of $d(A)$
- But perhaps within that a lot of variation is possible
- Joyce notes that there are some other constraints we can suggest flowing from the idea that we don't want to excessively reward either systematic moderation of immoderation

We won't, largely for time reasons, go over these constraints, save to note that they leave several rules open

Constraints on Scoring Rules

Joyce proves a remarkable result concerning the class of 'plausible' scoring rules

- If *Cr* is not a probability function, and *s* is a plausible scoring rule, there is some probability function *Pr* that is guaranteed to be more accurate than *Cr*
- That is, there is a probability function *Pr* whose accuracy is lower (i.e. better) than *Cr* however the world is
- So whatever way the world is, if your credences are *Cr*, you'll be more accurate if you change them to *Pr*

A Philosophical Consequence

This I think is a very nice result

- It does suggest that there's a purely doxastic reason to make credences be probabilities
- If credences are not probabilities, then you are introducing necessarily avoidable inaccuracy into your credal state
- And that seems not only bad, but bad in a way that suggests doxastic incoherence

Objection: No Numbers

Of course, this all assumes that credences are numerically representable

- We'll look presently at theories that reject this
- Joyce acknowledges that this is a major assumption
- He takes the result to be conditional; if all credences over propositions in $\mathcal{F}$ are numerical, they should form a probability function

Multiple Scoring Rules Objection

There's a tricky issue about the order of the quantifiers in Joyce's result that I don't quite know how to resolve. Joyce proved

- If Cr is not a probability function, then $\forall$ plausible scoring rules s, $\exists$ Pr that is, necessarily, more accurate than Cr according to s

He didn't prove

- If Cr is not a probability function, then $\exists$ Pr such that $\forall$ plausible scoring rules s, Pr is, necessarily, more accurate than Cr according to s

So it might be that an agent whose credences are Cr knows she can be more accurate, but doesn't know which Pr will make her more accurate, because she doesn't know the *right* scoring rule

Linear Scoring Rule Objection

Patrick Maher pointed out that the argument relies on one prima facie plausible scoring rule being ruled out

- That's the rule $s = d$

On that rule, there are plenty of Cr that are non-probabilistic, but such that we can't necessarily do better by moving to a probability function. Joyce has some arguments why that's a bad rule, but it is odd that this has to be ruled out.

Grounding Equal Probabilities

Consider the following two scenarios, concerning a (five-set) tennis match between Smith and Jones. (We'll say p = Smith wins.)

1. Your friend, who knows a bit about tennis, says that it should be a close match because the two players are pretty equally matched. At this point if you have to assign a credence to p, it arguably should be $\frac{1}{2}$

Grounding Equal Probabilities

Consider the following two scenarios, concerning a (five-set) tennis match between Smith and Jones. (We'll say p = Smith wins.)

1. Your friend, who knows a bit about tennis, says that it should be a close match because the two players are pretty equally matched. At this point if you have to assign a credence to p, it arguably should be $\frac{1}{2}$
2. After that, you watch the first four sets, which are split 2-2, with each set being very closely fought, and neither player having any obvious edge going into the fifth set. Again if you have to assign a credence to p, it arguably should be $\frac{1}{2}$, but you now arguably have more ground for this credence

Grounding Equal Probabilities

Consider the following two scenarios, concerning a (five-set) tennis match between Smith and Jones. (We'll say p = Smith wins.)

1. Your friend, who knows a bit about tennis, says that it should be a close match because the two players are pretty equally matched. At this point if you have to assign a credence to p, it arguably should be $\frac{1}{2}$
2. After that, you watch the first four sets, which are split 2-2, with each set being very closely fought, and neither player having any obvious edge going into the fifth set. Again if you have to assign a credence to p, it arguably should be $\frac{1}{2}$, but you now arguably have more ground for this credence

There's no simple way to represent the difference between your credal states in 1 and 2 in the Bayesian setting.

Low Credence in Exhaustive Propositions

If Cr is a probability function then either $Cr(A) \geq \frac{1}{2}$ or $Cr(\neg A) \geq \frac{1}{2}$. But perhaps this isn't always correct

- There are propositions about the distant past, or distant parts of space, where we don't have any grounds for having a high credence in either the proposition or its negation
- For a more every-day example, let $p =$ The price of gold in $US will be above the S&P500 at the end of 2038, and consider whether for you $Cr(p)$ is above or below $\frac{1}{2}$.

Low Credence in Exhaustive Propositions

If Cr is a probability function then either $Cr(A) \geq \frac{1}{2}$ or $Cr(\neg A) \geq \frac{1}{2}$. But perhaps this isn't always correct

- There are propositions about the distant past, or distant parts of space, where we don't have any grounds for having a high credence in either the proposition or its negation
- For a more every-day example, let $p =$ The price of gold in $US will be above the S&P500 at the end of 2038, and consider whether for you $Cr(p)$ is above or below $\frac{1}{2}$.

Low Credence in Exhaustive Propositions

If we want to say $Cr(A) < \frac{1}{2}$ and $Cr(\neg A) < \frac{1}{2}$, then we have to give up one of the following two principles

1. $Cr(A) + Cr(\neg A) = Cr(A \vee \neg A)$
2. $Cr(A \vee \neg A) = 1$

I'll go over the most prominent version of each option, then end with a more Bayesian friendly way of meeting these motivations.

Models for Shafer functions

For simplicity, assume the universe X is finite. As before, let $\mathcal{P}(X)$ be the powerset of X, and let m be a measure on $\mathcal{P}(X)$ such that $m() = 0$.

- Now consider a function $b : \mathcal{P}(X) \rightarrow [0, 1]$ defined as
- $b(X) = \sum Y \subseteq X m(Y)$
- That is, b is the sum of the measure of all the subsets of X

If all of the sets X such that $m(X) > 0$ are singletons, then b will be a probability function, but it won't be one in general.

Simple Shafer Functions

For simplicity, assume the universe X is finite. As before, let $\mathcal{P}(X)$ be the powerset of X, and let m be a measure on $\mathcal{P}(X)$ such that $m() = 0$.

- Now consider a function $b : \mathcal{P}(X) \rightarrow [0, 1]$ defined as
- $b(X) = \sum Y \subseteq X m(Y)$
- That is, b is the sum of the measure of all the subsets of X

If all of the sets X such that $m(X) > 0$ are singletons, then b will be a probability function, but it won't be one in general.

Example

Say we just care about a single proposition p

- So the universe basically consists of two possibilities, which we can call P and $\neg P$
- So $\mathcal{P}(\mathrm{X}) = \{\emptyset\{P\}, \{\neg P\}, \{P, \neg P\}\}$
- For a simple function, set $m(\{P\}) = m(\{\neg P\}) = m(\{P, \neg P\}) = \frac{1}{3}$
- Then $b(p) = b(\neg p) = \frac{1}{3}$
- But $b(p \vee \neg p) = 1$

Generality

Every probability function is a Shafer function, though not vice versa

- In fact, most of the alternatives to probability that have been proposed as quantitative theories of credence are kinds of Shafer functions
- So it's a nice very general approach to think about

Problems

We don't have nearly enough time to go over either of these problems, but there are three major reasons why the approach hasn't caught on

1. Hard to find a plausible approach to updating
2. Pearl's complaint: It confuses credence of p with credence of *Provably* p
3. More specific (in an odd way) than a non-numerical approach we'll soon consider

Drop LEM

If all you care about is that $Cr(p)$ and $Cr(\neg p)$ are less than $\frac{1}{2}$, there's an easy way to deal with this

- Drop the law of excluded middle!
- The natural way to do this is to base your theory around intuitionistic logic
- I'm the only person to have done this at more than a rudimentary level
- That is, my paper is the only one that considers this approach as a theory of credence

Intuitionist Probability Axioms

One nice approach about doing things this way is that we don't have to change the axioms at all

- We just interpret the entailment relation that is used in the axioms intuitionistically
- This impacts just about everything, since there are so many classical assumptions used from the start in probability theory
- But it is a natural starting point

Motivation

- Interestingly, many of the cases where theorists have wanted to say that $Cr(p)$ and $Cr(\neg p)$ are less than $\frac{1}{2}$ are cases where philosophers interested in realism have promoted kinds of anti-realism about p
- And if we're going to be anti-realists, we have independent reason to question LEM
- So this isn't entirely a technical fix

Kripke Models and Intuitionist Probability

Again, we won't go into this in any detail, but I'll just note that there is an easy to work with semantics for these probability functions

- Start with a normalised measure on a Kripke tree
- Say $Pr(A)$ is the measure of the set of points in the tree where A is forced
- Then the function will be an intuitionistic probability function

There is a lot of uncharted territory around here, since approximately every single theorem of the probability calculus that you see quoted relies on classical assumptions somewhere in its proof.

Motivating Intervals

In practice it is hard to figure out exactly what one's own credence in various propositions is

- It seems like this isn't just because of a failure of introspection
- Arguably there is no such thing as my credence that, say, a Democrat will win the next Presidential election
- And it isn't clear that this is a failure of rationality
- There isn't a reason that I should have a precise, numerical credence in this

Motivating Intervals

In practice it is hard to figure out exactly what one's own credence in various propositions is

- It seems like this isn't just because of a failure of introspection
- Arguably there is no such thing as my credence that, say, a Democrat will win the next Presidential election
- And it isn't clear that this is a failure of rationality
- There isn't a reason that I should have a precise, numerical credence in this
- Compare: It isn't a failure of rationality that I have no positive attitude towards some propositions p, neither believing them nor their negation

Intervals and Comparatives

This interacts interestingly with the representation theorem approach to constraints on credence

- One of the constraints we have to put in by hand is a kind of linearity
- $[Cr(A) \geq Cr(B)] \vee [Cr(B) \geq Cr(A)]$
- And that is hard to motivate
- Let A and B be propositions about separate future events that are both pretty likely, but it's hard to say exactly how likely
- It's not clear that one is obliged to have any comparative attitude about their likelihood

Blurry Probabilities

If you do that, you end up with the following nice system

- Every rational agent is representable by a **set** of probability functions S
- The agent's credence in A is higher than their credence in B iff for all $Pr \in S, Pr(A) > Pr(B)$
- And the agent's credence in A is at least as great as their credence in B iff for all $Pr \in S, Pr(A) \geq Pr(B)$

Radical Uncertainty

This leads to a nice possibility for representing agents who have, intuitively, no idea at all about p

- Just make sure that for every $x \in [0, 1]$ (or perhaps $(0, 1)$ there is a $Pr \in S$ such that $Pr(p) = x$

One really nice feature of this approach is that it allows this kind of uncertainty while retaining comparatives

- We can make the agent completely unsure, in this sense, about both p and q, while keeping it the case that (a) p has a higher credence than q, or (b) p and q are independent

Vague Interpretation

The definition of comparatives might remind one of supervaluational definitions of truth

- When I first started working on this approach, that's what I thought would be the right approach in general
- Each of the $Pr \in S$ is a precisification of the agent's credences, and any statement about the credences is true iff true on all precisifications
- That's actually how I started being interested in vagueness
- But now I think that's the wrong approach

Imprecise Interpretation

Consider again the idea from Joyce that we can evaluate credal states for accuracy

- When we evaluate non-probabilistic beliefs for accuracy, we have to judge accuracy against strength
- Consider, e.g. the following three people, with views about p - which is as it turns out true
  1. A believes p
  2. B believes $\neg p$
  3. C neither believes p nor believes $\neg p$
- There's a sense in which C's beliefs are as accurate as A's (at least no more inaccurate), but less strong

Imprecise Interpretation

I now think we should think of the agent that I referred to as having no idea about p as like C in the above example

- That's to say, her beliefs about p aren't inaccurate at all
- They simply aren't very strong
- And perhaps, given her evidence, that's an entirely rational response
- But note that on the supervaluationist move, there is no fact of the matter about how accurate her beliefs are
- That seems to me to be a mistake

Imprecise Interpretation

A few years ago, this sets of probability functions approach seemed like an interesting alternative to mainstream Bayesianism

- Now it seems to have basically taken over
- I still think there are a lot of interesting questions to address, e.g. about accuracy measurement, and about updating rules
- But it is now part of the mainstream
- And this, I think, is a little philosophical progress

The End