

Please turn in by 22 September 2009.

Required problems

1. Can your credence in a proposition that is compatible with your new information decrease when you update by conditionalization? In other words, can we have $f(q|p) < f(q)$ when p and q are not mutually exclusive? If so, give an example. If not, prove that this never happens.
2. Can your credence in a proposition that entails your new information decrease when you update by conditionalization? In other words, can we have $f(q|p) < f(q)$ when q entails p ? If so, give an example. If not, prove that this never happens.
3. Prove that conditionalizing on some new information preserves ratios of credences between propositions that entail that information. In other words, prove that when q and r each entail p ,

$$\frac{f(q|p)}{f(r|p)} = \frac{f(q)}{f(r)}$$

N.B. This result entails that conditionalizing on some evidence E preserves the ratio between your credence in a conjunction $(p \wedge E)$ and your credence in E , so preserves the conditional probability $f(p|E)$.

4. Derive $f(p \wedge q)$ as a function of $f(p)$ and $f(q)$, where f is a probability function with respect to which p and q are independent propositions.
5. In order to conditionalize your credence distribution on a proposition, your credence in that proposition must be well-defined. But suppose your credence distribution f is defined over a coarse algebra S , and you would like to update on some proposition that is not in that algebra. For instance: suppose you merely have credences about whether the temperature is in the 70's, 80's, or 90's, and you would like to update on the proposition that it is colder than 75 degrees.

Here is one rule for how to update your credences in this kind of situation: let A_1, A_2, \dots, A_n be the atoms of the algebra S , in the sense of JEFFREY 1965. For any atom A_i that is incompatible with your new information, define your updated credence function f' so that $f'(A_i) = 0$. Then renormalize your credence distribution over the remaining atoms, and define your credences in conjunctions of those atoms accordingly.

This strategy is in the spirit of updating by conditionalization, since it recommends that you confine your credence to those possibilities compatible with your evidence, and that you renormalize to make sure that your updated credence distribution is a probability measure. Explain why it is nevertheless a bad strategy for updating on a proposition that your prior credence distribution is not defined on.

6. "A jailer is in charge of three prisoners, called A, B, and C. They learn that two of the prisoners have been chosen at random to be executed the following morning; the jailer knows which two they are, but will not tell the prisoners. Prisoner A realizes that his chance of survival is $\frac{1}{3}$. But he says to the jailer, "Tell me the name of one of the other two who will be executed; this will not give me any information, because I know that at least one of them must die, and I know neither of them." The jailer agrees and tells A that B is going to be executed.

But now A has apparently raised his [credence that he will survive] from $\frac{1}{3}$ to $\frac{1}{2}$, because it will be either A or C who is the other one to be executed. Notice that if the jailer had said C instead of B, A's [credence that he will survive] would still apparently have gone up to $\frac{1}{2}$. How can this be?" (KELLY 1994)

Optional problems

7. State and prove the analog of the result in (3) for Jeffrey conditionalization. Then use the result to prove the characterization of Jeffrey conditionalization given in JEFFREY 1965:

$$\text{"(11-8): } \text{PROB}(A|A_i) = \text{prob}(A|A_i) \text{ for each } i = 1, 2, \dots, m.\text{" (174)}$$

8. JEFFREY 1965 begins with "the important special case in which $n = 2$, and where the pair B_1, B_2 has the form B, \bar{B} , so that we would be willing to describe the result of the observation simply by saying that it led the agent to change his degree of belief in some one proposition B " (168). Prove that any application of Jeffrey conditionalization can be factored into a finite sequence of changes to your credence in individual propositions, where each of those changes is an application of Jeffrey conditionalization on an assignment of credences to a partition with exactly two elements.

References

- JEFFREY, RICHARD C. 1965. "Probability Kinematics." In *The Logic of Decision*, 164–183. University of Chicago Press, Chicago.
- KELLY, D. G. 1994. *Introduction to Probability*. Macmillan, New York.

Please turn in by 6 October 2009.

Required problems

1. HOWSON & FRANKLIN 1994 mention that $I(P';P)$ is zero “if P' and P are the same distribution” (455). Demonstrate this result. N.B. Since $I(P';P)$ is non-negative, this entails that your current credence distribution is among the probability measures with the least divergence from itself. That’s a reassuring result; otherwise, the injunction to choose credences with least divergence from your current credences could mandate credence change even when no new constraints have been imposed on your credences.
2. Calculate $I(P';P)$ for the following probability functions:

$$P(x) = \frac{1}{3}, P(\bar{x}) = \frac{2}{3}$$

$$P'(x) = \frac{2}{3}, P'(\bar{x}) = \frac{1}{3}$$

3. Let x_1, x_2, x_3 be a partition of logical space. Calculate $I(P';P)$ for the following probability functions:

$$P(x_1) = \frac{1}{3}, P(x_2) = \frac{1}{2}, P(x_3) = \frac{1}{6}$$

$$P'(x_1) = \frac{2}{3}, P'(x_2) = \frac{1}{6}, P'(x_3) = \frac{1}{6}$$

Spell out how your answer is relevant to the following observation from HOWSON & FRANKLIN 1994: “if X and X' are two partitions, and X is finer than X' , then $I_{X'}(P';P) \leq I_X(P';P)$ ” (455).

4. WEINTRAUB 2004 describes the following variant on the original Sleeping Beauty case: “Sleeping Beauty is told she will see three lights flashing (one after the other), being made to forget what she has seen after each flash. If the (fair) coin lands heads, one of the three flashes will be red and two will be green. If the coin lands tails, one will be green and two will be red” (9).

WHITE 2006 describes the following generalized Sleeping Beauty case: “A random waking device has an adjustable chance $c \in (0,1]$ of waking Sleeping Beauty when activated on an occasion. In those circumstances in the original story where Beauty was awakened, we now suppose only that this waking device is activated. When $c = 1$, we have the original Sleeping Beauty problem. But if $c < 1$, the case is significantly different. For in this case Beauty cannot be sure in advance that she will be awakened at all during the experiment” (116).

Consider the following third variation on the Sleeping Beauty case, inspired by the above cases: Beauty is told that regardless of how a fair coin lands, she will see three lights flashing (one after the other), being made to forget what she has seen after each flash. If the coin lands heads, one of the three flashes will be red and two will be green. If the coin lands tails, a random device will be activated three times. The device has objective chance $c = \frac{1}{3}$ of flashing a red light and objective chance $c = \frac{2}{3}$ of flashing a green light.

Suppose Beauty sees a red flash in this third variation on the Sleeping Beauty case. What should be her credence that the fair coin landed heads? Justify your answer.

5. WHITE 2006 uses the generalized Sleeping Beauty case to present the following challenge for thirders: “if $c < 1$ then when Beauty wakes up she clearly does gain some information, namely

W : Beauty is awake at least once during the experiment.

And this is clearly relevant to whether

H : The coin landed Heads.

For the likelihood of W is greater given $\neg H$ than given H . Any answer must take into account the impact of this information on Beauty’s credence. But now the force of this impact must depend partly on the value of c . For the difference between the likelihoods $P_-(W|H)$ and $P_-(W|\neg H)$ increases as c decreases (where P_- is Beauty’s rational credence function prior to waking). The *degree* to which Beauty has a better chance of being awakened given two opportunities rather than one depends on how small c is. So whatever else we might say about Beauty’s rational credence in H when she wakes up, it should vary to some degree with the value of c . This is a result that the thirder, insofar as he follows the Elga and Arntzenius-Dorr arguments, cannot accommodate” (117-8).

Explain how this reasoning overgenerates to yield the incorrect verdict about what Beauty’s credence in heads should be when she wakes up in the third variation on the Sleeping Beauty case given above.

Optional problems

6. Consider the Judy Benjamin problem from VAN FRAASSEN 1989. Suppose Benjamin minimizes the relative information between her previous and updated credence distributions when she learns that she is .75 likely to be in the headquarters area if she is in Red Territory. What will be her updated credence that she is in Red Territory?
7. Suppose I know that the objective chance of rain tomorrow is .5. I falsely believe that it is .2 likely to flood if it rains. Suppose I then learn that there is no chance of flooding, since the city has piled sandbags high along the banks of every river.

Consider the following reasoning: since the proposition that it will flood entails the proposition that it will rain, the proposition that it will not flood confirms the proposition that it will not rain. But intuitively, despite learning that it will not flood, I should not lower my credence that it will rain. I should still think it is .5 likely to rain, and I should merely *locally* conditionalize on the proposition that it will almost certainly not flood. That is, I should lower my conditional credence that it will flood if it rains, from .2 to nearly 0, and I should re-normalize my credence in each maximally specific alternative in which it rains, so that the sum of my credence in these possibilities increases from .4 to nearly .5.

This reasoning appears to present a problem for the rule that we should update by conditionalization, since it appears to be a case where my credence in the hypothesis that it will rain should remain unchanged even when I learn a proposition that disconfirms that hypothesis. Explain why in fact there is no real problem for conditionalization here.

References

VAN FRAASSEN, BAS C. 1989. *Laws and Symmetry*. Oxford University Press, Oxford.

HOWSON, COLIN & ALLAN FRANKLIN. 1994. "Bayesian Conditionalization and Probability Kinematics." *British Journal for the Philosophy of Science*, vol. 45 (2): 451–466.

WEINTRAUB, RUTH. 2004. "Sleeping Beauty: A Simple Solution." *Analysis*, vol. 64: 8–10.

WHITE, ROGER. 2006. "The Generalized Sleeping Beauty Problem: A Challenge for Thirder." *Analysis*, vol. 66 (2): 114–19.

Please turn in by 20 October 2009.

Required problems

1. Consider the following sequence of utterances:

(1a) Tim expects that he will confuse his students.

(1b) So does Bob.

It is generally agreed among linguists that (1b) has three readings. Identify the contents of the three attitudes it can be used to ascribe, and say whether each is a *de se* or *de dicto* proposition.

2. Create a pair of bets that constitute a synchronic Dutch book for an agent who has .6 credence in a proposition and .6 credence in its negation. Hint: recall from the first handout that “Bayesians say that if $f(p)$ is your credence that p is true, then you should be willing to pay up to $f(p)$ dollars for a ticket that is worth one dollar if p is true and nothing if p is false.”
3. Suppose that you are currently certain that you are not the best singer in the world. But you think it is .5 likely that you are going to a karaoke bar tonight and that while under the influence of cheap beer and persuasive friends, you will be certain that you are the best singer in the world. Suppose a Dutch bookie offers you a bet where you pay \$20 if you are the world’s best singer and you go to the bar, you get \$20 if you are not the world’s best singer and you go to the bar, and you pay \$10 if you do not go to the bar. What further conditional bet can the bookie offer you in order to have constructed a diachronic Dutch book against you?
4. HALPERN 2006 aims to construct an argument against the thirder response to the original puzzle that is analogous to the pro-thirder argument given in HITCHCOCK 2004. Halpern says that if Beauty is a thirder, then she will accept bets in each of two betting situations:

“Before the experiment starts, Sleeping Beauty is offered a bet that pays off \$30 if the coin lands heads and \$0 otherwise, and costs \$15. Each time Sleeping Beauty is woken up, she is offered a bet that pays off \$30 if the coin lands tails and \$0 otherwise, and costs \$20, with the understanding that the bet pays off only once in each trial. In particular, if the coin in fact lands tails, and Sleeping Beauty takes the bet both times she is woken up, she gets the \$30 payoff only once (and, of course, only has to pay \$20 for the bet once). The accounting is done at the end of the trial. . . the second bet is fair to an agent who ascribes probability $2/3$ to tails when woken up.” (17)

Halpern argues that if Beauty accepts these bets, she is guaranteed to lose money. If the coin lands heads, she gains \$15 on the first bet but loses \$20 on the second. If the coin lands tails, she loses \$15 on the first bet and gains only \$10 on the second. Halpern concludes that Beauty should not have $1/3$ credence in heads upon waking.

Consider the following variation on the Sleeping Beauty case: suppose experimenters put Beauty to sleep and flip a fair coin on Sunday. If the coin lands heads, they wake Beauty once on Monday by ringing a bell and once on Tuesday by sounding an alarm. If the coin lands tails, they wake her once on Monday and once on Tuesday, ringing a bell both times. Regardless, they administer a drug on Monday night to make her forget her Monday waking. Beauty knows all this.

Intuitively, what should Beauty's credence in heads be when she wakes up to the sound of a bell ringing? Demonstrate that Halpern's reasoning in the passage quoted above leads to an unacceptable conclusion in this variation on the Sleeping Beauty case.

5. Give an example of two representors that yield the same vague credence for p and the same vague credence for q , while yielding different vague credences for $p \wedge q$.

Optional problem

6. Give an example of two representors that yield the same vague credence for p and the same vague credence for q and the same vague credence for $p \vee q$, while yielding different vague credences for $p \wedge q$.

References

HALPERN, JOSEPH Y. 2006. "Sleeping Beauty Reconsidered: Conditioning and Reflection in Asynchronous Systems." In *Oxford Studies in Epistemology*, TAMAR SZABÓ GENDLER & JOHN HAWTHORNE, editors, vol. 1, 111–42. Oxford University Press, Oxford.

HITCHCOCK, CHRISTOPHER READ. 2004. "Beauty and the Bets." *Synthese*, vol. 139 (3): 405–420.

Please turn in by 3 November 2009.

Required problems

1. State a critical question for Adam Elga's paper on imprecise credences.
2. JOYCE 2009 presents the following case: suppose I have a fair coin that can land heads or tails, and two other coins that can land black or white, and suppose that "you know the black/white coins are *complementary* – the first's bias toward heads is β while the second's is $(1 - \beta)$ – but you have no other information about β 's value. For all you know, β might be any number in $[0, 1]$. I will toss the fair coin. If a head comes up, I'll toss the first black/white coin. If a tail comes up, I'll toss the second black/white coin. Before the coin is tossed, how confident should you be of black? After you observe a head/tail, how confident should you be of black?" Justify your answers to both questions.
3. Suppose that an agent has only finitely many credence distributions in her representor. Recall from seminar that if she acts according to the arithmetic mean of those credence distributions, then we will be able to construct a diachronic Dutch book against her. Prove this claim by demonstrating that conditionalizing credence distributions on a given proposition does not commute with taking the arithmetic mean of those credence distributions, where the arithmetic mean h of probability functions f and g is defined as $h(p) = \frac{f(p) + g(p)}{2}$.
4. Prove that conditionalizing credence distributions on a given proposition *does* commute with taking the geometric mean of those credence distributions, where the geometric mean h' of f and g is defined as $h'(p) = \sqrt{f(p) \cdot g(p)}$.
5. Does taking the arithmetic mean of two credence distributions preserve unanimous judgments to the effect that two given propositions are independent?
6. Let p and q be propositions, and let S be the set of probability functions that give a certain conditional credence to q given p . Give examples demonstrating that the following two operations on a representor R sometimes but not always coincide:
 - a. intersecting R with a set of probability functions S
 - b. using probability kinematics to update each member of R on the constraint corresponding to S
7. In VAN FRAASSEN 2006, van Fraassen constructs the following example: Peter has high credence that he will get a good grade on a certain test. "He has at the outset no opinion at all about whether it will be cloudy tomorrow, the day of the new test. Nor is there anything in his opinion that bears on whether the weather and his grades are connected or not connected" (486). Then van Fraassen points out the following: when Peter wakes up and sees that it is cloudy, he will have no opinion about whether he will get a good grade, i.e. his credence that he will get a good grade on the test will dilate completely.

According to van Fraassen, this example presents “devastation” for those who accept imprecise credences. He notes that Peter’s credences will dilate and says:

How could this be? The reason was not that he had a tacit conditional expectation of meteorological influences on his test performance! On the contrary, the reason was precisely that he *had no opinion at all* about the presence or absence of any weather-grade correlation. Vagueness on any such unrelated topic will therefore unsettle and destabilize the content of his opinion when information on that topic comes in... Updating on apparently *irrelevant* bits of news can seemingly be destructive of one’s painfully acquired legitimate expectations. (488, 484)

Explain why van Fraassen’s example does not actually demonstrate that fans of imprecise credences must accept the disturbing conclusion that people should ordinarily go around becoming less confident of propositions about test results simply because they observe certain weather events.

Optional problem

8. A set S of probability functions is *convex* just in case for any f and g in S , the weighted arithmetic mean $h_a(p) = a \cdot f(p) + (1 - a) \cdot g(p)$ is also in S , for any real number $a \in [0, 1]$. Demonstrate that a set of probability functions according to which two given propositions are independent is not always convex. Explain how your example is relevant to the first section of VAN FRAASSEN 1990, in which van Fraassen formally characterizes various probabilistic opinions.

References

- VAN FRAASSEN, BAS C. 1990. “Figures in a Probability Landscape.” In *Truth or Consequences: Essays in Honor of Nuel Belnap*, J.M. DUNN & A. GUPTA, editors, 345–356. Dordrecht.
- . 2006. “Vague Expectation Value Loss.” *Philosophical Studies*, vol. 127: 483–491.
- JOYCE, JAMES M. 2009. “Do Imprecise Credences Make Sense?” Prague Foundations of Uncertainty Conference, September 4, 2009.

Please turn in by 24 November 2009.

Required problems

1. If you update by conditionalization, the ratio of your updated credences in the hypotheses H_1 and H_2 should equal the product of the ratio of your prior credences in those hypotheses and the *likelihood ratio*, which is defined as follows:

$$\frac{P(E|H_1)}{P(E|H_2)}$$

Demonstrate this result by proving the following formula:

$$\frac{P(H_1|E)}{P(H_2|E)} = \frac{P(H_1)}{P(H_2)} \cdot \frac{P(E|H_1)}{P(E|H_2)}$$

2. Some argue that likelihood ratios play a crucial role in defining how disagreeing agents should compromise. Read the following discussion by economist Robin Hanson:

<http://robinhanson.typepad.com/overcomingbias/2009/02/share-likelihood-ratios-not-posterior-beliefs.html#more>

Apply the compromise technique covered in Hanson's discussion to the following simple case. Suppose that Jane and James are trying to determine whether a particular urn contains one black ball and one white ball, or two white balls. Jane and James each start with .5 credence in those two hypotheses. Jane privately draws two balls from the urn, both of which are white. James privately draws one ball from the urn, which is white.

Use likelihood ratios to calculate the credence that Jane and James each should have in the two urn hypotheses after they draw balls from the urn privately. Calculate the joint credence they would have in the hypotheses if they averaged their credence distributions. Calculate the joint credence they would have in the hypotheses if they updated by trading likelihood ratios. Which compromise method yields the more intuitive result?

3. Suppose I know that the objective chance of rain tomorrow is .5. I falsely believe that it is .2 likely to flood if it rains. Suppose I then learn that there is no chance of flooding, since the city has piled sandbags high along the banks of every river.

Consider the following reasoning: since the proposition that it will flood entails the proposition that it will rain, the proposition that it will not flood confirms the proposition that it will not rain. But intuitively, despite learning that it will not flood, I should not lower my credence that it will rain. Rather, I should still think it is .5 likely to rain, and I should merely *locally* conditionalize on the proposition that it will almost certainly not flood. That is, I should lower my conditional credence that it will flood if it rains, from .2 to nearly 0, and I should re-normalize my credence in each maximally specific alternative in which it rains, so that the sum of my credence in these possibilities increases from .4 to nearly .5.

This reasoning appears to present a problem for the rule that we should update by conditionalization, since it appears to be a case where my credence in the hypothesis that it will rain should remain unchanged even when I learn a proposition that disconfirms that hypothesis. Explain why in fact there is no real problem for conditionalization here.

Hint: according to conditionalization, you should update by conditionalizing your current credences on the strongest proposition that you learn. In this example, you learn something stronger than that there is no chance of flooding. Be sure to take proper account of this fact by representing your earlier belief state with a credence distribution over a sufficiently fine-grained algebra.

4. There are several ways to extend the theory of updating in Moss 2009 to the standard Sleeping Beauty case. Suppose that when Beauty wakes up, her sense of time passing is indifferent between two alternatives: that one day has passed since Sunday and that two days have passed. Determine what credence Beauty should then have in the proposition that the coin landed heads.

Suppose instead that when Beauty wakes up, her sense of time passing tells her that if the coin landed heads, then it is Monday, whereas if the coin landed tails, it is .5 likely that one day has passed and .5 likely that two days have passed. Determine what credence Beauty should then have in the proposition that the coin landed heads.

5. As mentioned in seminar, JOYCE 2005 has some odd consequences for cases where an agent conditionalizes on the information that a certain objective chance hypothesis does not obtain. Construct an example in which an agent comes to rule out the proposition that the objective chance of some hypothesis X is a certain value, and the quantity $w(X, E)$ thereby decreases, where E is her total evidence.

Joyce says “the weightier the evidence for X is, the smaller $w(X, E)$ will tend to be” (166) and also that “the size or ‘weight’ of the evidence has to do with how much relevant information the data contains” (159). These remarks seem to entail that the agent in the case you constructed does not gain information relevant to the hypothesis X . Is there an intuitive defense of this result?

References

- JOYCE, JAMES M. 2005. “How Probabilities Reflect Evidence.” *Philosophical Perspectives*, vol. 19: 153–178.
- MOSS, SARAH. 2009. “Updating as Communication.” Ms., Department of Philosophy, University of Michigan.