Do Imprecise Credences Make Sense?

Jim Joyce
Department of Philosophy
The University of Michigan
jjoyce@umich.edu

Many people – including Issac Levi, Peter Walley, Teddy Seidenfeld, Richard Jeffrey and Mark Kaplan – have suggested that uncertain beliefs in light of evidence are best represented by *sets* of probability functions rather than individual probability functions. I will defend the use of such "imprecise credal states" in modeling beliefs against some recent objections, raised by Roger White and others, which concern the phenomenon of dilation. Dilation occurs when learning some definite fact forces a person's beliefs about an event to shift from a fully determinate subjective probability to an imprecise spread of probabilities. A number of commentators have found aspects of dilation disturbing, both from an epistemic and decision-theoretic perspective, and have placed the blame on the idea that beliefs can be imprecise. I shall argue that these worries are based on an overly narrow conception of imprecise belief states which assumes that we know everything there is to know a person's doxastic attitude toward an event once we know all possible values of her subjective probability for that event. A well-developed theory of imprecise beliefs has the resources to characterize a rich family of relationships among doxastic attitudes that are essential to a complete understanding of rational belief. It can only do so, however, if it is purged of an overly narrow conception of belief states. Once this change is made dilation does not seem so disturbing, and certain decision problems that often seem perplexing can be resolved in a satisfactory manner.

The Basic Ideas of Bayesian Epistemology

- Believing is not an all-or-nothing matter. Opinions come in varying gradations of strength ranging from full certainty of truth to complete certainty of falsehood.
- Gradational belief is governed by the laws of probability, which codify the minimum standards of consistency (or "coherence") to which rational opinions must confirm.
- Learning involves *Bayesian conditioning*: a person who acquires data D should modify her opinions in such a way that her "posterior" views about the relative odds of propositions consistent with D agree with her "prior" views about these relative odds.
- Gradational beliefs are often revealed in decisions. Rational agents choose options they *estimate* will produce desirable outcomes, and these estimates are a function of their beliefs.

	X & Y	X & ~Y	~X & Y	~X & ~Y
Option - O_X	prize	prize	penalty	penalty
Option - O_Y	prize	penalty	prize	penalty

You should (determinately) prefer O_X to O_Y if and only if you are more confident of X than of Y.

"Precise" Bayesianism

- Graded beliefs come in sharp numerical degrees: in any context, we can think of an agent's epistemic state as given by a single *credence function* $\boldsymbol{c}$ that assigns a degree of belief $\boldsymbol{c}(X) \in [0, 1]$ to each proposition X (in some Boolean algebra).

  Note: "Objective" and "subjective" Bayesians differ about whether credence functions reflect objective constraints on beliefs or (also) matters of personal opinion.

- Rational credences are *additive*: $\boldsymbol{c}(X \vee Y) = \boldsymbol{c}(X) + \boldsymbol{c}(Y)$ when X and Y are contraries.
- Learning is governed by *Bayes' Theorem*: a person who acquires data D should modify her opinions so that she believes each proposition X to degree

  $$\boldsymbol{c}^*(X) = \boldsymbol{c}(X)\cdot[\boldsymbol{c}(D \mid X)/\boldsymbol{c}(D)]$$

  where $\boldsymbol{c}(D \mid X) = \boldsymbol{c}(D \mathbin{\&} X) / \boldsymbol{c}(X)$ is the prior conditional probability of D given X.
- Rational decision making is a matter of choosing options that maximize *expected utility* computed relative to one's credences.

Problems with Precise Degrees of Belief

- It is psychologically unrealistic to suppose that people have attitudes that are precise enough to be represented by real numbers. What could $\boldsymbol{c}(X) = 1/\pi$ mean?
- Since evidence is often incomplete, imprecise or equivocal, the *right* response is often to have beliefs that are incomplete, imprecise or equivocal.

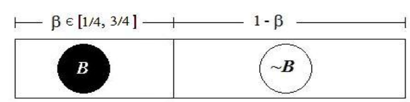

A black/white coin is chosen randomly from an urn containing coins of every possible bias ¼ < β < ¾. You have *no information* about the proportions with which coins of various biases appear in the urn.

How confident should you be that the coin comes up black when next tossed?

- "Objective" Bayesian: $\boldsymbol{c}(B) = ½$ because this choice uniquely minimizes the amount of *extra information* one needs to add to get a sharp degree of belief.
- "Imprecise" Bayesian: It is determinate that $¼ < \boldsymbol{c}(B) < ¾$, but $\boldsymbol{c}(B)$ lacks a determinate value because the evidence is insufficiently precise to discriminate among assignments $\boldsymbol{c}(B) = p$ with $¼ < p < ¾$.

- The Precise Interpretation misrepresents *uncertainty* in decision making.

♁ THE IMPRECISE INTERPRETATION ♁

- IMPRECISION. Graded beliefs do not come in sharp degrees. A person's *credal state* is typically best represented by a family C of degree-of-belief functions.
- COHERENCE. For a rational agent, C will contain only probability functions.
- CONDITIONING. If a person with credal state C learns that some proposition D is certainly true, then her post-learning credal state will be $C_D = \{c(\bullet \mid D) : c \in C\}$.
- SUPERVALUATION. Truths about what a person believes correspond to properties shared by *every* credence function in the person's credal state.
  (Converse? Part of our topic today.)

Important qualification: imprecise ≠ vague, ≠ unknown.

OMISSION

I will say little about how C is determined in light of evidence.

Nearly all my examples will involve two coins: one that comes up heads/tails and is fair, and one that comes up black/white and whose bias is unknown to some degree.

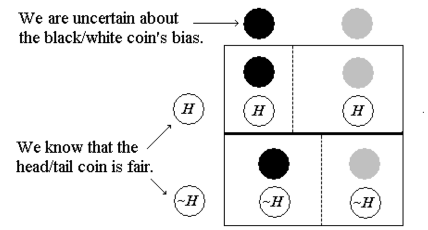

I will assume that the credal state of a person who knows only that the bias β of the black/white coin toward *Black* is between p and q is a set of credence functions C such that $p < c(Black) < q$ for all $c \in C$.

DIFFERENCE FROM SHARP VIEWS

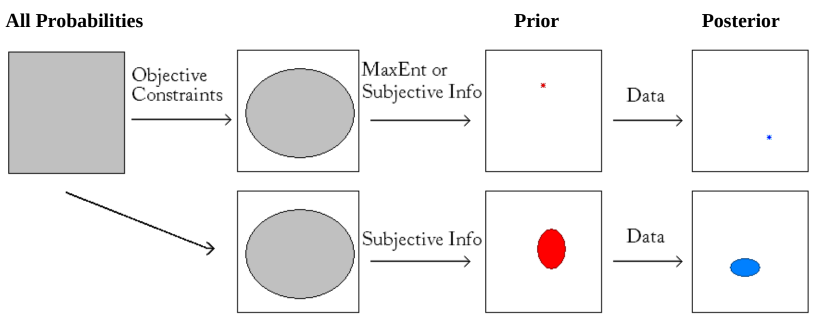

Symmetry Condition (precise): If $E_1,\ldots, E_N$ are "equipossible" given the constraints, then the prior should assign them the same probability $c(E_i) = c(E_j)$.

Symmetry Condition (imprecise): If $E_1,\ldots, E_N$ are "equipossible" given the constraints, then for any credal state $c \in C$ and any permutation σ of $\{1, 2,\ldots, N\}$ there exists a $c_\sigma \in C$ such that $c_\sigma(E_i) = c(E_{\sigma(i)})$.

A KEY QUESTION ABOUT THE IMPRECISE MODEL

Which of features of the model reflect facts about what the person believes, and which are mere artifacts of the formalism?

- Clearly not all features of C reflect facts about the person's beliefs. E.g., each credence function assigns a sharp probability to every event, but the person doesn't.

Two Common (but I say wrong) Answers:

1. LOWER ENVELOPE. If C and C^* assign the same lower (and so upper) probabilities to all propositions in some Boolean algebra $\mathcal{A}$ of propositions, then C and C^* describe the same beliefs about the propositions in $\mathcal{A}$.

   - The lower probability that C assigns to X is $C_-(X) = \inf\{c(X): c \in C\}$; the upper probability is $C^+(X) = \sup\{c(X): c \in C\}$; coherence ensures $C_-(X) = 1 - C^+(\sim X)$.
   - I think of $C_-(X)$ as the lowest probability that the evidence for X allows, and of $C^+(X)$ is the highest probability for X that the evidence against X (= for $\sim X$) allows.

Example

You have been exposed to a dread disease. Between 30% and 60% of people exposed contract the virulent form (the "+" condition). The rest contract a tamer version (the "–" condition). The virulent form kills 70% if those who have it, while the tamer form kills only 10%.

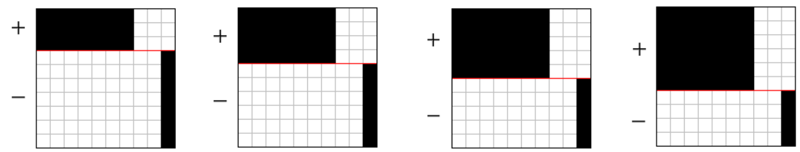

$c(+) = 0.3$ $c(D) = 0.28$ | $c(+) = 0.4$ $c(D) = 0.34$ | $c(+) = 0.5$ $c(D) = 0.4$ | $c(+) = 0.6$ $c(D) = 0.46$

According to LOWER, all the useful information in this picture is at the two ends. When we know that the lower probability of death is $\mathcal{C}_{-}(D) = 0.28$ and the lower probability of life is $\mathcal{C}_{-}(\sim D) = 0.54$, we know all there is to know.

2. SET. If $\mathcal{C}$ and $\mathcal{C}^*$ assign the same ranges to all propositions in $\mathcal{A}$, so that $\mathcal{C}(X) = \mathcal{C}^*(X)$ for all $X \in \mathcal{A}$, then $\mathcal{C}$ and $\mathcal{C}^*$ describe the same beliefs about the propositions in $\mathcal{A}$.

Notation/Terminology: Let $\mathcal{C}(X) = \{\boldsymbol{c}(X): \boldsymbol{c} \in \mathcal{C}\}$ denote the set of all values assigned to the proposition X by credence functions in $\mathcal{C}$, the $\mathcal{C}$-range of X.

- SET and LE agree when $\mathcal{C}$ is a convex set since in that case $\mathcal{C}(X)$ and $\mathcal{C}^*(X)$ will be intervals with common upper and lower bounds.

My Aim: To explain why SET/ LOWER fails and thereby to refute a number of recent objections to imprecise probabilities that presuppose them.

A Consequence of Set/Lower

Imagine a black/white coin and a red/blue coin. Nothing is known about their biases. Consider these credal states where you know nothing about the bias of the black/grey coin or about the bias of the red/blue coin:

Independence

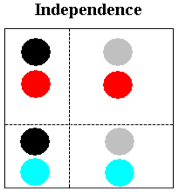

$\boldsymbol{c}(B \mid R) = \boldsymbol{c}(B \mid \sim R), \forall \boldsymbol{c} \in C$

Complementarity

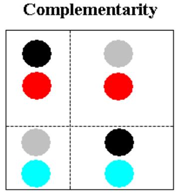

$\boldsymbol{c}(B \mid R) = 1 - \boldsymbol{c}(B \mid \sim R), \forall \boldsymbol{c} \in C$

Since all propositions in the algebra generated by $\{B, R\}$ have lower probabilities of zero (except $\top$) in each state, SET and LOWER **wrongly** deem these credal states equivalent.

- In Independence, learning R or $\sim R$ makes no difference to one's beliefs about B and $\sim B$.
- In Complementarity, one's beliefs about B upon learning R are determinately *contrary* to one's beliefs about B upon learning $\sim R$.

Why Lower/Set Fails

According to SET, we know everything there is to know a person's doxastic attitudes in re propositions in $\mathcal{A}$ once we have identified all sets $\mathcal{C}(X) = \{\boldsymbol{c}(X): \boldsymbol{c} \in \mathcal{C}\}$ for $X \in \mathcal{A}$. This ignores relationships among such sets that clearly do reflect facts about beliefs.

Example-1 (Dominance): Suppose the black/white coin has bias β toward black, and the red/blue coin has bias of $2\beta - {}^4/_3\beta^2$ toward red, where, for all you know, β can have any value strictly between 0 and ¾. $\mathcal{C}(B) = \mathcal{C}(R) = (0, ¾)$, yet you are determinately more confident in R than in B since $\boldsymbol{c}(R) = 2{\cdot}\boldsymbol{c}(B) - {}^4/_3{\cdot}\boldsymbol{c}(B)^2 > \boldsymbol{c}(B)$ for all $\boldsymbol{c} \in \mathcal{C}$.

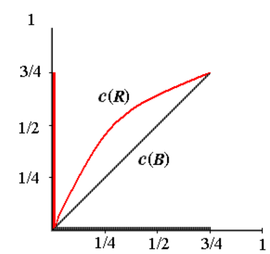

Example-2 (Complementarity): Let the black/white coin's bias be β and the red/blue coin's bias be 1 – β, where it is only known that β falls between ¼ and ¾. Again, $\mathcal{C}(B) = \mathcal{C}(R)$. But, since the coins are anti-correlated, black and red are *complementary*, i.e., $\boldsymbol{c}(B) = 1 - \boldsymbol{c}(R)$ for all $\boldsymbol{c} \in \mathcal{C}$.

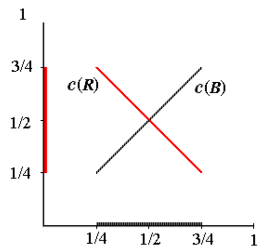

Useful Representation (Directed Intervals): $\mathcal{C}(R) = [0, 1]$ and $\mathcal{C}(B) = [1, 0]$.

Note: Directions are *relative*, the direction of $\mathcal{C}(R)$ is defined relative to that of $\mathcal{C}(B)$.

Example-3 (Conditional Complementarity): The red/blue coin's bias is β if the black/white coin lands black and 1 – β if it lands white, where it is only known that β falls between ¼ and ¾. Consistent with your beliefs, red can have any probability between ¼ and ¾ conditional on black or on white. But, you have complementary attitudes toward the red outcome depending on what you learns about the black/white coin: $\boldsymbol{c}(R \mid B) + \boldsymbol{c}(R \mid {\sim}B) = 1$ for all $\boldsymbol{c} \in \mathcal{C}$.

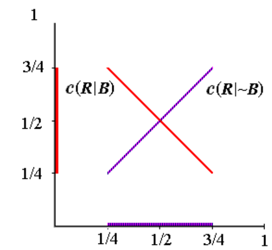

The Puzzling Phenomenon of Dilation

I have a *fair* heads/tails coin, and two black/white coins which are *complementary* – the first's bias is β while the second's is (1 – β). You have no information about β's value.

I will toss the fair coin. If a head comes up, I'll toss the first *B*/*W* coin. If a tail comes up I'll toss the second *B*/*W* coin.

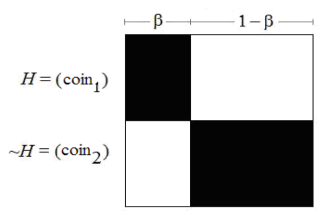

For all you know, β might be any number in [0, 1].

- 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*?

The Right Answer

(Seidenfeld and Wasserman, 1993)

- Your credal state $\mathcal{C}$ contains every probability with $\boldsymbol{c}(H) = ½$ and $\boldsymbol{c}(B \mid H) = \boldsymbol{c}({\sim}B \mid {\sim}H)$.
- Even though you know zilch about β, your prior credence for *Black* is *determinately* ½.
- After you learn something very specific your credal state will become $\mathcal{C}_H$ or $\mathcal{C}_{\sim H}$.
- Either way, you credence for *Black* will cover the whole range [0, 1].

How should we understand this puzzling change?

- It appears that, whatever you learn about the fair coin, you are certain to *lose* a great deal of information about the *Black*/*White* coin.
- It appears that your sharp credence of $\boldsymbol{c}(B) = ½$ is unwarranted since you know that you won't have evidence to back it up *when you know more*.

Two Objections to Credal States that Admit of Dilation

- *"Narrowing" Objection*. Such credal states violate:

  Narrowing. Learning the truth-value of a proposition should always *narrow* the permissible credal states since adding information excludes possibilities.

  In dilation the set of permissible credences for *B* expands from {½} to [0, 1], and we *lose* information about *B* upon learning *H* or ~*H* (Halpern and Grünwald, 2004).

- *"Reflection" Objection*. Such credal states violate:

  Weak Reflection. If you know you will undergo a veridical learning experience in which you will discover *H*'s truth-value (and nothing else), and if you can now identify a specific credal state for *B* that you will inhabit *whatever* you learn, then you should now be inhabiting that credal state right now.

  (White, 2009): Before the fair coin is tossed you know you will be maximally uncertain about *B* whatever you learn. So, you should be maximally uncertain now!

The Common Mistake

Both objections assume, with Set, that you end up in the *same* belief state relative to {*Black*, ~*Black*} whichever way the fair coin falls. But…

- Before you learn how the fair coin falls you see the beliefs you will have about *B* if you learn *H* as *complementary* to the beliefs you will have about *B* if you learn ~*H*.
  - If you learn *H*, you will come to know that *B* and ~*B* are perfectly correlated with *B* & *H* and ~*B* & *H*, respectively.
  - If you learn ~*H*, you will come to know that *B* and ~*B* are perfectly correlated with *B* & ~*H* and ~*B* & ~*H*, respectively.
  - Your imprecise prior credences for *B* & *H* and ~*B* & *H* are *complementary* in the sense that $\boldsymbol{c}(B \,\&\, H) + \boldsymbol{c}({\sim}B \,\&\, H) = ½$ for all $\boldsymbol{c} \in C$. And, *B* & ~*H* and ~*B* & ~*H* are imprecise and complementary in the same way.
  - Critically, both $\boldsymbol{c}(B \,\&\, H) = \boldsymbol{c}(B \,\&\, {\sim}H)$ and $\boldsymbol{c}({\sim}B \,\&\, H) = \boldsymbol{c}({\sim}B \,\&\, {\sim}H)$ hold for all $\boldsymbol{c} \in C$.

Your sharp ½ credence for *B* thus results from your symmetrical uncertainty about which imprecise, complementary pair — ⟨*B* & *H*, ~*B* & *H*⟩ or ⟨*B* & ~*H*, ~*B* & ~*H*⟩ — is perfectly correlated with ⟨*B*, ~*B*⟩.

Response to Narrowing Objection

While the Narrowing Principle is correct, the objection wrongly assumes that you lose information about *B* when you learn *H* or ~*H*.

In fact, you gain information about *B*!

- You learn which of the two correlations obtains.
- Since each of these correlations precludes the other, the new evidence provided by learning *H* or ~*H* removes your uncertainty about which complementary pair — ⟨*B* & *H*, ~*B* & *H*⟩ or ⟨*B* & ~*H*, ~*B* & ~*H*⟩ — is correlated with ⟨*B*, ~*B*⟩.
- This new evidence undermines your rationale for $\boldsymbol{c}(B) = ½$.

Response to Reflection Objection

Weak Reflection is right, but the objection wrongly assumes that you think you will end up in the *same* imprecise posterior credal state regarding *B* whatever you learn about *H*.

- The objector says you end up in state $C_H(B) = [0, 1]$ if *H* and in state $C_{\sim H}(B) = [0, 1]$ if ~*H*. According to Lower Envelope or Set, these states are identical.

In fact, these states are *conditionally complementary* from the perspective of your prior.

- As we have seen, Lower Envelope and Set ignore the fact that, the identity $\boldsymbol{c}(B \mid H) = 1 - \boldsymbol{c}(B \mid {\sim}H)$ holds relative to every credence function in your prior.
- Thus, you **do not** know in advance what your post-toss beliefs will be! You do not know whether you should align your future beliefs about ⟨*B*, ~*B*⟩ with your current beliefs about ⟨*B* & *H*, ~*B* & *H*⟩ or with your complementary current beliefs about ⟨*B* & ~*H*, ~*B* & ~*H*⟩.

In light of this, the only way to satisfy WR is by having a credal state in which *B*'s current probability is determinately identical to its expected future probability, which is exactly what you do!

WHITE'S OBJECTION

White (2009) proposes a counterexample to the idea of imprecise credal states:

- I will toss a fair heads/tails coin and (independently) toss a black/white coin about which you know nothing. I see both results; you see neither.
- If I see a head, I'll report the true color of *B*/*W*. If a tail, I'll report the wrong color.
- In effect, you learn $H \equiv B$ (if I say "black") or $H \equiv \sim B$ (if I say "white"), where "≡" is the material conditional. Since the head/tail coin is fair and tosses are independent, you anticipate these with equal probability: $\boldsymbol{c}(H \equiv B) = \boldsymbol{c}(H \equiv \sim B) = ½$ for all $\boldsymbol{c} \in C$.

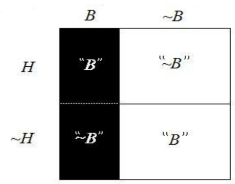

Before hearing what I say, what should you believe about *B*?

Once I say "black" ($H \equiv B$), what should you believe about *H* and *B*?

WHITE'S ANSWER: You should believe *H* and *B* to degree ½ (both before and after the fair coin is tossed)!

- If *B*'s prior is imprecise over [0, 1] then when I say "black" (or "white"), *H*'s posterior probability should coincide with *H*'s known prior objective chance of ½.
  - Hence, *B*'s posterior probability must also be determinately ½ (since you will then know $B \equiv H$ and $\boldsymbol{c}(B \mid H \equiv B) = \boldsymbol{c}(H \mid H \equiv B)$).
  - Hence, *B*'s **prior** probability must also be determinately ½**!**

    Note: White does not assert this, but he must accept it if the posterior is to be the result of conditioning on the data since, via the independence of the tosses, $\boldsymbol{c}(B \mid H \equiv B) = \boldsymbol{c}(B)$ for all $\boldsymbol{c} \in C$.
  - *Reductio*: If *B*'s prior is maximally imprecise, then it is precisely ½**.**

TITELBAUM'S (2009) "TECHNICOLOR AL" (retold)

Al is playing tag tomorrow. Either he or Hal will be **It**. Before bed Al learns he is **It**, but Al won't remember in the morning. The rules stipulate that a referee will privately toss a fair coin, and will give the person who is **It** a blue hat if heads and a red hat if tails.

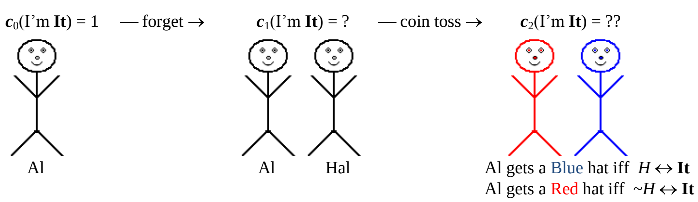

Titelbaum:

- $\boldsymbol{c}_0(H) = \boldsymbol{c}_1(H) = \boldsymbol{c}_2(H) = ½$ (because the coin is known to be fair throughout).

  $\boldsymbol{c}_2$(I'm **It**) $= \boldsymbol{c}_2(H) = ½$ when $\boldsymbol{c}_2$(*Blue*) = 1 and $\boldsymbol{c}_2$(I'm **It**) $= \boldsymbol{c}_2(\sim H) = ½$ when $\boldsymbol{c}_2$(*Red*) = 1. So, whatever color Al sees, $\boldsymbol{c}_2$(I'm **It**) = ½.

  Since $\boldsymbol{c}_2$ must arise from $\boldsymbol{c}_1$ by conditioning, $\boldsymbol{c}_1$(I'm **It** | *Blue*) = $\boldsymbol{c}_1$(I'm **It** | *Red*) = ½.
- So, we get $\boldsymbol{c}_1$(I'm **IT**) = ½ ≠ [0, 1] *without invoking the Principle of Indifference*.

THE KEY CLAIM

Both arguments require that learning of a correlation between the fair head/tail coin and the uncertain black/white coin provides no evidence about the outcome of the fair toss.

Put differently: Information about the correlation is *admissible* with respect to the chance statement "the head/tail coin is fair."

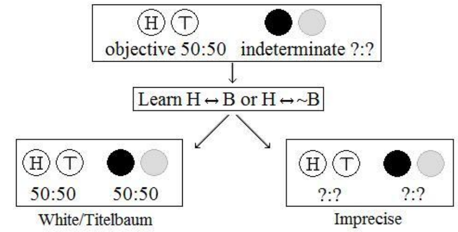

The laws of probability require the values in the bottom box to be the same, but do not say whether $\boldsymbol{c}(H) = \boldsymbol{c}(B) = ½$ or whether *H* and *B* might have the same *imprecise* value. (What does "same imprecise value" mean? Good question!)

White argues for the Key Claim as follows:

- You started out knowing head/tail was fair.
- You lost no information, nor have you acquired inadmissible information that undermines your knowledge of the chances.
- Of course, if you had independent reason to think B true (or false) then learning of a correlation between H and B *would* be relevant to your beliefs in H.
- But… "*you haven't a clue whether as to whether B*." So, when you learn $H \equiv B$, "you have nothing to suggest that the coin landed heads or that it landed tails. Shouldn't you just ignore this useless bit of information and keep your credence in heads at ½" thereby keeping your credences aligned with known objective chances?

☆ Question: Is $H \equiv B$ evidentially relevant to H when one knows nothing about B?

When one learns of a perfect correlation between an event of known objective chance and an event about which one knows nothing, should one retain one's beliefs about the chance event, even if this means changing one's views about the unknown event?

Is $H \equiv B$ evidentially relevant to H?
Five Arguments for Thinking Not, With Refutations

1. (Silly) Since the sharp credence $\boldsymbol{c}(H) = ½$ is based on known chances it can't ever be undermined by the discovery of a correlation between H and any event whose chance is imprecisely known.

❖ If we knew that the black/white coin was biased between 0.95 and 0.99 toward heads, we would surely conclude that the probability of heads was in that range too.

2. (Less silly) Since $\boldsymbol{c}(H) = ½$ is based on known chances it can't ever be undermined by the discovery of a correlation between H and an event whose chance is imprecisely known and whose uncertainty is symmetrical around ½.

❖ Why does the symmetry matter? Presumably, because it ensures that discovering the correlation is evidentially equivalent to discovering the reverse correlation. But is this actually so? See next argument.

Likelihood Principle

3. $H \equiv B$ can be evidence for or against H only if, from the perspective of the prior, there is a difference between it's probability conditional on H and its probability conditional on $\sim H$. But, there is no such difference since $\boldsymbol{c}(H \equiv B \mid H)$ and $\boldsymbol{c}(H \equiv B \mid \sim H)$ can have any value in [0, 1], consistent with the prior.

❖ The first sentence is true, but the second fails to recognize that $H \equiv B$ given H is the complement of $H \equiv B$ given $\sim H$.

Conditioning on H changes the probability of $H \equiv B$ to twice the (indeterminate) probability of H & B, whereas conditioning on $\sim H$ changes the probability of $H \equiv B$ to twice the (indeterminate) probability of $\sim H$ & B.

These are determinately different probabilities since H & B and $\sim H$ & B are complementary according to $\boldsymbol{c}$.

Shouldn't $H \equiv B$ be Evidentially Relevant B?

4. Since $\boldsymbol{c}(B \mid H \equiv B) = \boldsymbol{c}(B)$, $\forall \boldsymbol{c} \in \mathcal{C}$, the proposed account makes $H \equiv B$ *evidentially irrelevant* to B. This seems wrong: learning that B is perfectly correlated with H should change one's evidential situation with respect to B (as well as H).

   - Worse, even if we imagine a head/tail coin with any bias $\alpha \in (0, 1)$, learning $H \equiv B$ still has no impact on B's credence, which remains entirely imprecise.

     But, it seems clear that as α moves higher/lower learning about the correlation should raise/lower B's probability more and more drastically.

❖ Fortunately, learning $H \equiv B$ *does* alter B's (imprecise) probability in the right way.

- Increasing the bias toward heads increases in B's (imprecise) probability, *in the dominance sense.*

  When $x > y$, $\boldsymbol{c}(B \mid H \equiv B \text{ and } \alpha = x) > \boldsymbol{c}(B \mid H \equiv B \text{ and } \alpha = y)$, $\forall \boldsymbol{c} \in \mathcal{C}$.

- Moreover, the effects of reversing the bias of heads/tails affect the probability of B *in the complementarity sense.*

  For $0 < x < 1$, $\boldsymbol{c}(B \mid H \equiv B \text{ and } \alpha = x) > 1 - \boldsymbol{c}(B \mid H \equiv \sim B \text{ and } \alpha = x)$, $\forall \boldsymbol{c} \in \mathcal{C}$.

An Important Moral About Evidential Relevance

The framework of imprecise probability offers many ways of representing evidential relationships. It is a mistake to interpret $\boldsymbol{c}(B \mid H \equiv B) = \boldsymbol{c}(B)$, $\forall \boldsymbol{c} \in \mathcal{C}$, as meaning that learning $H \equiv B$ leaves the evidential status B unchanged.

- In one sense, learning the truth about $H \equiv B$ is evidentially irrelevant to B.

  After all, the resolution of the $H \equiv B$ vs. $H \equiv \sim B$ question does not affect B's probability. The balance of evidence for B is the same whatever one learns.

- In another sense, the truth about $H \equiv B$ is highly evidentially relevant to B.

  In the event that $H \equiv B$, the probability of B is directly related to that of H and inversely related to that of $\sim H$.

  In the event that $H \equiv \sim B$, the probability of B is inversely related to that of H and directly related to that of $\sim H$.

  What one learns about $H \equiv B$ alters the "directional valence" of the evidential dependence of B on H.

Where is the Inadmissible Information?

5. To undermine knowledge of objective chances one should be able to point to some "inadmissible" information. Where is it?

   Definition: Relative to a credence function $\boldsymbol{c}$, E is inadmissible with respect to X when $\boldsymbol{c}(X \mid \text{chance}(X) = x \,\&\, E) \neq x$.

❖ As White himself notes, learning $H \equiv B$ undermines our knowledge of the head/tail coin's fairness relative to every credence function in $\mathcal{C}$ with $\boldsymbol{c}(B) \neq ½$. Specifically,

$$\boldsymbol{c}(H \mid \beta = ½ \,\&\, H \equiv B) > ½ \text{ when } \boldsymbol{c}(B) > ½$$

$$\boldsymbol{c}(H \mid \beta = ½ \,\&\, H \equiv B) < ½ \text{ when } \boldsymbol{c}(B) < ½$$

So, since B's probability is maximally indeterminate, it is maximally indeterminate whether, how much, and in which direction, one's knowledge of H's chance will be undermined by learning the correlation.

Given this "indeterminate inadmissibility" one's credences about H *should* become indeterminate when the correlation between H and B is learned.

Some Useful References

Christensen, D. 2004. *Putting Logic in its Place: Formal Constraints on Rational Belief*. Oxford: Clarendon Press.

Couso, I., Moral, S. and Walley, P. 2000. "A survey of concepts of independence for imprecise probabilities," *Risk, Decision and Policy*: 165-181.

Gaifman, H. 1985. "A Theory of Higher Order Probabilities." In *Causation, Chance, and Credence*, ed. Brian Skyrms and William Harper, 191–219. Boston: Kluwer Academic.

Grünwald, P. and Halpern, J. 2004. "When Ignorance is Bliss," *ACM International Conference Proceeding Series; Vol. 70: Proceedings of the 20th conference on Uncertainty in artificial intelligence*.

Grove, A.J. and Halpern, J. 1998. 'Updating Sets of Probabilities,' *Proceedings of the Fourteenth Conference on Uncertainty in AI*, 173–182.

Joyce. J. "How Degrees of Belief Reflect Evidence," *Philosophical Perspectives* **19**, 2005: 153-179.

Levi, I. 1974. "On Indeterminate Probabilities", *Journal of Philosophy* **71**: 397-418.

Levi, I. 1980. *The Enterprise of Knowledge*, Cambridge: M.I.T. Press.

Seidenfeld, T. and Wasserman, L. 1993. "Dilation for Convex Sets of Probabilities," *Annals of Statistics* **21**: 1139–1154.

Titelbaum, M. "Feigning Indifference," forthcoming

van Fraassen, B. 1984. "Belief and the Will," *The Journal of Philosophy* **81**:

van Fraassen, B. 2006. 'Vague Expectation Loss' *Philosophical Studies* **127**: 483-491.

Walley, P. 1991. *Statistical Reasoning with Imprecise Probabilities*, London: Chapman and Hall.

White, R. "Evidential Symmetry and Mushy Credence," forthcoming in Oxford Studies in Epistemology.