Thoughts on Decision Making with Imprecise Probabilities

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

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*.

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.

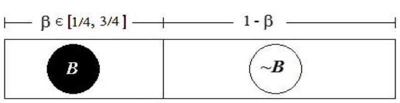

A black/white coin is chosen randomly from an urn containing coins of every possible bias $\frac{1}{4} < \beta < \frac{3}{4}$. You have *no information* about the proportions with which coins of various biases appear in the urn.

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

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

Illustration: The Four Urns

A rational person's beliefs will "reflect her evidence," but beliefs reflect evidence in a variety of ways, not all of which are adequately captured by sharp probabilities.

Imagine that coin drawn from the following urns will be tossed.

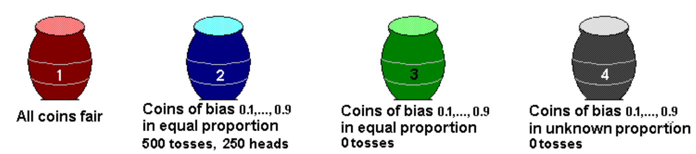

- *Popper's Objection*: Bayesianism treats all four cases as identical by assigning $\boldsymbol{C}(Heads) = \frac{1}{2}$, but these are entirely different evidential situations.
- *Jeffrey's Reply*: same credence ≠ same epistemic state.
  - The Urn_1 and Urn_2 probabilities are *resilient*. They remain fixed (exactly for Urn_1, roughly for Urn_2) given future evidence: $\boldsymbol{c}(Heads \mid 24\text{H}, 1\text{T}) = \frac{1}{2}$ and $\boldsymbol{c}(Heads \mid 24\text{H}, 1\text{T}) \approx 0.50023$.
  - The Urn_3 probabilities are *unstable* in the face of evidence: $\boldsymbol{c}(Heads \mid 24\text{H}, 1\text{T}) = 0.880$.

What about Urn_4? What value should we assign $\boldsymbol{c}(Heads \mid 24\text{H}, 1\text{T})$?

"OBJECTIVE" BAYESIAN ANSWER

Objective Bayesian: Urn_4 and Urn_3 are equivalent. ***c***(*Heads* | *Data*) exists and equals ***c***(*Heads* | *Data*) for all data.

- Key Claim: In both cases, no evidence distinguishes one bias (among {0.1,…, 0.9}) from any other. Principles of sound epistemology (Insufficient Reason, MaxEnt) require that we treat symmetrical cases symmetrically *by assigning them the same probability*.

  The choice of ***c***(*Heads* | *data*) = ***c***(*Heads* | *Data*) is often justified by appeal to the requirement that the prior probability ***c*** should encode the minimum amount of information consistent with the evidence, so that $\boldsymbol{c}(\beta = 0.i) = p_i$ maximizes $\text{Entropy}(p_i) = -\sum_i p_i \cdot ln(p_i)$.

  Problem(s): ''Apart from evolving a vitally important piece of knowledge, that of the exact form of the distribution, out of an assumption of complete ignorance, it is not even a unique solution.'' R.A. Fisher, 1922, pp. 324-325.

"SUBJECTIVE" BAYESIAN ANSWER

Subjective Bayesian: Urn_4 and Urn_3 are not equivalent. ***c***(*Heads* | *Data*) exists but it can consistently have any value in [0, 1], whatever the value of ***c***(*Heads* | *Data*).

- Key Claim: A rational agent can have any probabilistically coherent set of credences over the possible biases. So, any credence for heads in light of data can be rationally entertained.
- Problem: The choice of any sharp probability over any other seems arbitrary. In particular, in the face of symmetrical evidence there is no more reason to choose a prior with $\boldsymbol{c}(\beta = 0.i) = p_i$ than the symmetrical prior with one with $\boldsymbol{c}(\beta = 1 - 0.i) = p_i$.

A THIRD WAY

"The problem is not that Bayesians have yet to discover the truly noninformative priors, but rather that no precise probability distribution can adequately represent ignorance."

(P. Walley, *Statistical Reasoning with Imprecise Probabilities*, 1991)

Imprecise Bayesian: Urn_4 and Urn_3 are not equivalent. ***c***(*Heads* | *Data*) does not even exist! Instead of trying to model our beliefs about possible biases using a single credence function, we should use a ***set*** of credence functions that best reflects our true state of uncertainty.

- *Objective Imprecise*: There is a single imprecise credal state that is appropriate for any given body of evidence. For symmetric evidence this state is symmetric.
- *Subjective Imprecise*: There are typically many *sets* of credence functions consistent with the data, and a believer is free to adopt any of these as her credal state.

Judgments in light of evidence are seen as imposing qualitative "constraints" on credal states**.** E.g.,

X is more likely that *Y*	$\boldsymbol{c}(X) > \boldsymbol{c}(Y)$ for all $\boldsymbol{c} \in C$
X and *Y* are independent	$\boldsymbol{c}(X \,\&\, Y) = \boldsymbol{c}(X) \cdot \boldsymbol{c}(Y)$ for all $\boldsymbol{c} \in C$
X and *Y* are (+, –) correlated	$\boldsymbol{c}(X \,\&\, Y) >,< \boldsymbol{c}(X) \cdot \boldsymbol{c}(Y)$ for all $\boldsymbol{c} \in C$

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 $C_D = \{\boldsymbol{c}(\bullet \mid D) : \boldsymbol{c} \in C.\}$.
- SUPERVALUATION. Truths about what a person believes correspond to properties shared by *every* credence function in the person's credal state.

White: Imprecise Probabilities Lead to Absurd Decisions

Roger White asks: If your credence in a proposition is imprecise over an interval $[x, y]$, at what odds should you bet? He considers two "common" answers.

Liberal. It is mandatory to take bets that maximize expected utility according to all $\boldsymbol{c} \in C$, and it is permissible to take bets that maximize expected utility according to some $\boldsymbol{c} \in C$.

Conservative. It is mandatory to take bets that maximize expected utility according to all $\boldsymbol{c} \in C$, and it is impermissible to take any bets other than these.

- I'll focus on Liberal, but it will become clear that both it and Conservative are poor ways to make decisions with imprecise probabilities.

NOTE: I think that the categorical notions of a *permissible* act or an *impermissible* act are entirely out of place in an imprecise framework.

White's Decision Problem

I'll independently toss a *fair* head/tail coin and a black/white coin of entirely unknown bias. You don't see the outcome, but I will report either $H \equiv B$ or $H \equiv {\sim}B$.

After my report, you must accept or decline this *unfavorable* wager:

BET = [\$1 if H, –\$2 if $\sim H$] (objective expected payoff = –\$0.5)

Note: If your credence for H were determinately greater than $^2/_3$ you'd accept BET, and if it were determinately less than $^2/_3$ you'd reject BET.

- After my report the range of your permissible credences for H comes to cover the entire interval $[0, 1]$ *whether you learn* $H \equiv B$ or $H \equiv {\sim}B$.
- As a result, your utility for BET after my report is imprecise over the interval $[–2, 1]$.

You should reject BET before you learn what I say, but what should you do *after*?

Doxastic Imprecision = Practical Incoherence?

White: If you are a liberal, it is *permissible*, for purposes of action, for you to assign H any probability whatever after learning $H \equiv B$ or $H \equiv {\sim}B$. In particular, you can assign H a probability that exceeds $^2/_3$ once you hear my report.

- So, says White, it is permissible for you to accept BET after my announcement, no matter what I announce!
- You thereby ensure yourself of *an objective expected loss* of 50¢.
- White: Liberals may permissibly accept BET *before* the announcement.

That's nuts! And if imprecise probabilities required this it would be a serious problem.

Doxastic Imprecision ≠ Practical Incoherence

White's liberal uses the following strategy:

- *Crude*. Select an action-guiding credence function such that $\boldsymbol{c}(H \mid A) > {}^2/_3$ whatever announcement $A \in \{H \equiv B, H \equiv {\sim}B\}$ is made.

$$C(H) = \{1/2\} \xrightarrow[\text{}]{\text{I announce } H\equiv B \text{ or } H \equiv \sim B} C(H) = [0, 1] \xrightarrow{\text{You choose for action}} C(H) = ({}^2/_3, 1] \longrightarrow \text{Accept}$$

White's agent is using a decision rule

- When $C(H) = [0, 1]$ act as if $C(H) = ({}^2/_3, 1]$.

This makes sense only if your post-announcement epistemic state, $C(H) = [0, 1]$, is the same *whatever you learn*. But it's not the same!

From the perspective of your initial credal state, the epistemic states you might be in after hearing my announcement are *conditionally complementary*:

$$\mathbf{c}(H \mid H \equiv B) + \mathbf{c}(H \mid H \equiv {\sim}B) = \mathbf{c}(B) + \mathbf{c}({\sim}B) = 1 \text{ for all } \mathbf{c} \in C.$$

So, even though your credence in heads will cover the whole of the interval [0, 1] whatever I learn, your imprecise beliefs about H in the two situations are contrary.

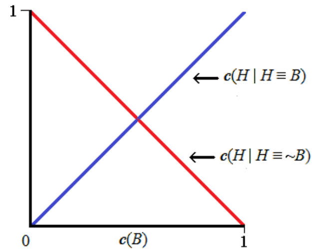

It follows that the utilities for BET are complementary too:

$$\text{EU}_{\mathbf{c}}(\text{BET} \mid H \equiv B) + \text{EU}_{\mathbf{c}}(\text{BET} \mid H \equiv {\sim}B) = -0.5 \text{ for all } \mathbf{c} \in C.$$

DOXASTIC IMPRECISION ≠ PRACTICAL INCOHERENCE

Recognizing this complementarity, a more sophisticated liberal would use:

Sophisticated Liberal. Select an action-guiding credence function with $\mathbf{c}(H \mid A) > {}^2/_3$ if one announcement is made, and select an action-guiding credence with $\mathbf{c}(H \mid A) \leq {}^2/_3$ if the other announcement is made.

I announce $H \equiv B$ / You choose for action:

$C(H) = \{1/2\}$ ⟶ $C(H) = [0, 1]$ ⟶ $C(h) = ({}^2/_3, 1]$ ⟶ Accept

$H \equiv {\sim}B$:

$C(H) = \{1/2\}$ ⟶ $C(H) = [1, 0]$ ⟶ $C(h) = [0, {}^2/_3]$ ⟶ Decline

White's agent uses the *information suppressing* rule:

- When $C(H) = [0, 1]$ act as if $C(h) = ({}^2/_3, 1]$; when $C(H) = [1, 0]$ act as if $C(h) = ({}^2/_3, 1]$.

Instead a liberal should use:

- When $C(H) = [0, 1]$ act as if $C(h) = ({}^2/_3, 1]$; when $C(H) = [1, 0]$ act as if $C(h) = [0, {}^2/_3]$.

A GOOD WAY OF PUTTING THE POINT (DUE TO ADAM ELGA)

"Your state of imprecision makes you imprecise as to the evidential relevance of $H \equiv B$ for H. You are imprecise over a range of answers to the question: is '$H \equiv B$ evidence for or against H?' But every answer to that question that endorses taking the bet upon learning $H \equiv B$ will endorse rejecting the bet upon learning $H \equiv {\sim}B$. That is why you shouldn't be poised to accept BET in response to both pieces of news."

BIGGER WORRY: CAN A RATIONAL STRATEGY *EVER* ACCEPT BET?

Even if a sophisticated liberal is not required to accept BET whatever she learns, it still seems wrong that it should *ever* be permissible to accept. After all, as White points out, the bet is objectively disadvantageous.

> If we repeat the betting scenario many times a person who, say, accepts BET if $H \equiv B$ and rejects it if $H \equiv {\sim}B$, is almost certain to lose lots of money.

The problem is that Liberal (and Conservative) are lousy decision rules in the context of imprecise credences because they fail to recognize *symmetric complementarities*.

	H & ($B \equiv H$)	~H & ($B \equiv H$)	H & ($B \equiv {\sim}H$)	~H & ($B \equiv {\sim}H$)
$\text{Bet}_{B \equiv H}$	\$1	– \$2	\$0	\$0
$\text{Bet}_{B \equiv {\sim}H}$	\$0	\$0	\$1	– \$2

Note, by way of foreshadowing, how $\text{Bet}_{B \equiv {\sim}H}$ is obtained from $\text{Bet}_{B \equiv H}$ by exchanging payoffs in complementary cases (color coded).

Imprecise Decision Making

An imprecise decision rule is a mapping from credal states into *partial* choice functions.

- For each credal state C and set of actions $\mathcal{A}$, a partial choice function will return two (possibly empty) subsets of $\mathcal{A}$
  - $\text{Choice}_c(\mathcal{A}) \subseteq \mathcal{A}$ is the set of acts that are determinately preferred to all others, and so determinately *choiceworthy*.
  - $\text{Reject}_c(\mathcal{A}) \subseteq \mathcal{A}$ is the set of acts that are determinately dispreferred to all others, and so determinately *prohibited*.

(Note: A fuller theory will *rank* acts by their levels of choiceworthiness.)

Some Imprecise Truths about Rational Decision Making
(in the simple case where we have a *single* utility function)

Consistency. $\text{Choice}_c(\mathcal{A})$ and $\text{Reject}_c(\mathcal{A})$ are disjoint.

Dominance. Let $\text{EU}_c(a)$ be the expected utility of act a computed using the credence function $c \in C$, then

- If $\text{EU}_c(a) \geq \text{EU}_c(b)$ for all $c \in C$ and all $b \in \mathcal{A}$, then $a \in \text{Choice}_c(\mathcal{A})$.
- If $\text{EU}_c(a) < \text{EU}_c(b)$ for all $c \in C$ and all $b \in \mathcal{A}$, then $a \in \text{Reject}_c(\mathcal{A})$.

In other words, if a determinately maximizes expected utility, then it is choiceworthy. If it determinately does not maximize expected utility, then it is should be rejected.

Independence. If $[a \text{ if } E, d \text{ if } {\sim}E] \in \text{Choice}_c\{[a \text{ if } E, d \text{ if } {\sim}E], [b \text{ if } E, d \text{ if } {\sim}E]\}$ for some act d, then this holds for every act d. Same for Reject.

Irrelevant Alter. If $A \in \text{Choice}_c\mathcal{A}$ and $B \in \text{Reject}_c\mathcal{A} \cup \{B\}$ then $A \in \text{Choice}_c\mathcal{A} \cup \{B\}$. If $A \in \text{Reject}_c\mathcal{A}$ and $B \in \text{Choice}_c\mathcal{A} \cup \{B\}$ then $A \in \text{Reject}_c\mathcal{A} \cup \{B\}$.

Three Grades of Coherence

Suppose, for all $c \in C$, $\text{EU}_c(a) = \text{EU}_c(a_1) + \text{EU}_c(a_2)$ and $\text{EU}_c(b) = \text{EU}_c(b_1) + \text{EU}_c(b_2)$.

- From C's perspective, a is a "package" of a_1 and a_2, b is a "package" of b_1 and b_2.

Basic Idea of Coherence: your attitudes toward the elements of a package should be consistent with your attitudes about the package as a whole.

Coherence (weak): If a is judged choiceworthy against b, then a_1 and a_2 cannot both be rejected in comparison with b_1 and b_2, respectively. If b is rejected in comparison with a, then b_1 and b_2 cannot both be choiceworthy against a_1 and a_2, respectively.

$$a \in \text{Choice}_c\{a, b\} \Rightarrow a_1 \notin \text{Reject}_c\{a_1, b_1\} \text{ or } a_2 \notin \text{Reject}_c\{a_2, b_2\}$$

$$b \in \text{Reject}_c\{a, b\} \Rightarrow b_1 \notin \text{Choice}_c\{a_1, b_1\} \text{ or } b_2 \notin \text{Choice}_c\{a_2, b_2\}.$$

Coherence (mid): If a is judged choiceworthy against b, then one part of a can only be rejected in favor of its b-counterpart if the other part of a is judged choiceworthy against its b counterpart. If b is rejected in comparison with a, then one part of b can only be choiceworthy against its a-counterpart if the other part of b is rejected in comparison with its a counterpart.

$$a \in \text{Choice}_c\{a, b\} \Rightarrow a_1 \in \text{Reject}_c\{a_1, b_1\} \text{ only if } a_2 \in \text{Choice}_c\{a_2, b_2\}$$

$$b \in \text{Reject}_c\{a, b\} \Rightarrow b_1 \in \text{Choice}_c\{a_1, b_1\} \text{ only if } b_2 \in \text{Reject}_c\{a_2, b_2\}.$$

Coherence (strong): If a is judged choiceworthy against b, then at least one part of a must be judged choiceworthy against its b counterpart. If b is rejected in comparison with a, then at least one part of b must be rejected in comparison with its a counterpart.

$$a \in \text{Choice}_c\{a, b\} \Rightarrow a_1 \in \text{Choice}_c\{a_1, b_1\} \text{ or } a_2 \in \text{Choice}_c\{a_2, b_2\}$$

$$b \in \text{Reject}_c\{a, b\} \Rightarrow b_1 \in \text{Reject}_c\{a_1, b_1\} \text{ or } b_2 \in \text{Reject}_c\{a_2, b_2\}.$$

Assessment

At the moment, I am inclined to regard Coherence-mid as plausible, and Coherence-strong as too strong (but my views on these issues are unstable).

Here is something that might look like a counterexample to Coherence-strong:

	X	$\sim X$
a_1	10	0
b_1	1	8
a_2	0	10
b_2	8	1

where every value for $\boldsymbol{c}(X)$ in (0, 1) can be found in your credal state.

It is tempting to say that (i) you should have no determinate preference between a_1 and b_1 since some of your credal states favor one and some favor the other, (ii) same for a_2 and b_2, but (iii) dominance requires you to prefer a_1 & a_2 to b_1 & b_2.

On the other hand, this ignores the fact that (i*) any credal state that favors b_1 over a_1 will also favor a_2 over b_2 by a larger margin, (ii*) the same is true with "1" and "2" reversed, and (iii*) credal states with ${}^8/_{17} < \boldsymbol{c}(X) < {}^9/_{17}$ favor both a_1 and a_s.

A Further Requirement: Symmetric Ignorance → Same Choice

Definition: Say that act a^μ is the ***mirror*** of act a just in case

- a^μ and a are identical except that a produces payoff u in event E and payoff of v in event E^* whereas a^μ produces payoff of v in E and payoff of u in E^*.
- E and E^* are complementary, so that $\boldsymbol{c}(E) + \boldsymbol{c}(E^*) = k$ for all $\boldsymbol{c} \in C$.
- E and E^* are *evidentially symmetric* in the sense that for every $x \in [0, k]$, if $\boldsymbol{c}(E) = x$ for some $\boldsymbol{c} \in C$ then there is a $\boldsymbol{c^*} \in C$ with $\boldsymbol{c^*}(E^*) = x$.

Example of Mirroring:

You know only that E's objective chance is between ${}^1/_4$ and ${}^3/_4$.

	E	$\sim E$
Wager on E	\$1	\$0
Wager on $\sim E$	\$0	\$1

Parity

The relationship between mirrors is symmetric like indifference, but strictly weaker. Call its transitivity closure *parity* (= "being on a par").

Parity. If a and a^* are on a par, then for any set of actions $\{b_1,\ldots, b_n\}$

$a \in \text{Choice}_c\{a, b_1,\ldots, b_n\}$ if and only if $a^* \in \text{Choice}_c\{a^*, b_1,\ldots, b_n\}$

$a \in \text{Reject}_c\{a, b_1,\ldots, b_n\}$ if and only if $a^* \in \text{Reject}_c\{a^*, b_1,\ldots, b_n\}$

Parity requires that you treat "Wager on E" and "Wager on $\sim E$" the same way when choosing among options. When one is choiceworthy so is the other. When one is rejected so is the other. When one is neither so is the other.

NOTE!! This does ***not*** require you to be *indifferent* between the wagers. Indeed, it is reasonable to have

$\text{Choice}_c\{\text{Wager on } E, \text{Wager on } \sim E\} = \text{Reject}_c\{\text{Wager on } E, \text{Wager on } \sim E\} = \varnothing$

An Example of Parity:

	H & $(B \equiv H)$	$\sim H$ & $(B \equiv H)$	H & $(B \equiv \sim H)$	$\sim H$ & $(B \equiv \sim H)$
$\text{BET}_{B \equiv H}$	\$1	– \$2	\$0	\$0
M	\$0	– \$2	\$1	\$0
$\text{BET}_{B \equiv \sim H}$	\$0	\$0	\$1	– \$2
Z	\$0	\$0	\$0	\$0

Red events are complementary. Blue events are complementary.

The act M is the symmetric mirror of both $\text{BET}_{B \equiv H}$ and $\text{BET}_{B \equiv \sim H}$, which means that the latter two acts are on a par.

So, according to Parity, $\text{BET}_{B \equiv H}$ and $\text{BET}_{B \equiv \sim H}$ must play exactly the same role in your decision making.

In particular, if one of them is choiceworthy (rejected) against Z then the other must be choiceworthy (rejected) against Z as well.

COMBINING PARITY AND COHERENCE

Since

- Dominance ensures that $Z \in \text{Choice}_c\{\text{BET}, Z\}$ and $\text{BET} \in \text{Reject}_c\{\text{BET}, Z\}$.
- For all $\boldsymbol{c} \in C$, $EU_c(\text{BET}) = EU_c(\text{BET}_{B \equiv H}) + EU_c(\text{BET}_{B \equiv \sim H})$
  $EU_c(Z) = EU_c(Z) + EU_c(Z)$.

we can combine Parity and Coherence to obtain the following results:

Parity + Coherence-mid $\Rightarrow$

$\text{BET}_{B \equiv H} \notin \text{Choice}_c\{\text{BET}_{B \equiv H}, Z\}$ and $\text{BET}_{B \equiv \sim H} \notin \text{Choice}_c\{\text{BET}_{B \equiv \sim H}, Z\}$

$Z \notin \text{Reject}_c\{\text{BET}_{B \equiv H}, Z\}$ and $Z \notin \text{Reject}_c\{\text{BET}_{B \equiv \sim H}, Z\}$

That is, it is determinate that neither component of BET is choiceworthy, and that Z cannot be rejected in a comparison with either component.

Parity + Coherence-strong $\Rightarrow$

$\text{BET}_{B \equiv H} \in \text{Reject}_c\{\text{BET}_{B \equiv H}, Z\}$ and $\text{BET}_{B \equiv \sim H} \in \text{Reject}_c\{\text{BET}_{B \equiv \sim H}, Z\}$

$Z \in \text{Choice}_c\{\text{BET}_{B \equiv H}, Z\}$ and $Z \in \text{Choice}_c\{\text{BET}_{B \equiv \sim H}, Z\}$

That is, it is determinate that each component of BET should be rejected in a comparison with Z, and that Z should be chosen.

Note that this is stronger than the previous in saying that you should definitely reject each component of BET, not merely that you would be wrong to choose it.

CONDITIONING: A PRINCIPLE OF DYNAMIC CHOICE

Conditioning. If $\boldsymbol{C}^*$ is the credal state that you will have upon learning E, then

$[a \text{ if } E, d \text{ if } \sim E] \in \text{Choice}_c\{[a \text{ if } E, d \text{ if } \sim E], [b \text{ if } E, d \text{ if } \sim E]\} \Rightarrow a \in \text{Choice}_{c^*}\{a, b\}$.

$[a \text{ if } E, d \text{ if } \sim E] \in \text{Reject}_c\{[a \text{ if } E, d \text{ if } \sim E], [b \text{ if } E, d \text{ if } \sim E]\} \Rightarrow a \in \text{Reject}_{c^*}\{a, b\}$.

This says that your posterior preferences upon learning E are determined by your prior preferences conditional on E.

Consequences:

- Given Coherence-Mid: Upon learning either $B \equiv H$ or $B \equiv \sim H$, accepting BET will not be determinately choiceworthy and turning it down (i.e., accepting Z) will not be determinately prohibited.
- Given Coherence-Strong: Upon learning either $B \equiv H$ or $B \equiv \sim H$, accepting BET will be determinately prohibited, and turning it down (i.e., accepting Z) will be determinately choiceworthy.

ANSWERING WHITE'S OBJECTION, 2 WAYS

Proponents of Coherence-Mid will say:

- It is enough that BET is determinately not choiceworthy, and that Z is determinately not prohibited, after $B \equiv H$ or $B \equiv \sim H$ is learned.
- Decision theory should not tell us more than this: whether you learn $B \equiv H$ or $B \equiv \sim H$, choosing Z is never determinately worse than choosing BET and choosing BET is never determinately better than choosing Z.

Proponents of Coherence-Strong will say:

- It is not enough! Rationality determinately prohibits the choice of both $\text{BET}_{B \equiv H}$ and $\text{BET}_{B \equiv \sim H}$ before anything is learned, and this means that choosing BET is remains prohibited whether $B \equiv H$ or $B \equiv \sim H$ is learned.
- So, White is *wrong*. Bayesianism with imprecise probabilities is perfectly capable of ruling out the irrational choice of BET, provided that one subscribes to the principles of Parity, Conditioning, and Coherence-strong.

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.