“Frank Agnosticism”

“The rules for belief functions permit us, when we have little evidence bearing on a proposition, to express frank agnosticism by according both that proposition and its negation very low degrees of belief. [...]

“The Bayesian theory, on the other hand, cannot deal so readily with the representation of ignorance, and it has often been criticized on this account. The basic difficulty is that the theory cannot distinguish between lack of belief and disbelief.”

- Shafer 1976 *A Mathematical Theory of Evidence*, pp.22-23.

Belief Functions: The Axioms

A *belief function* on a finite set W is a function Bel: $2^W \rightarrow [0,1]$ satisfying:

cf. probability measures, which need only be defined on a subalgebra of 2^W

1. $\mathrm{Bel}(\emptyset) = 0$

2. $\mathrm{Bel}(W) = 1$

3. $\mathrm{Bel}(U \cup V) \geq \mathrm{Bel}(U) + \mathrm{Bel}(V) - \mathrm{Bel}(U \cap V)$

$\mathrm{Bel}(U \cup V \cup W) \geq \mathrm{Bel}(U) + \mathrm{Bel}(V) + \mathrm{Bel}(W) - \mathrm{Bel}(U \cap V) - \mathrm{Bel}(U \cap W) - \mathrm{Bel}(V \cap W) + \mathrm{Bel}(U \cap V \cap W)$

would be "=" for probability measures

etc.

"Frank Agnosticism" from the Axioms

Take an event and its complement: A, $\neg A$.

The three axioms imply:

$$1 = \mathrm{Bel}(W) = \mathrm{Bel}(A \cup \neg A) \geq \mathrm{Bel}(A) + \mathrm{Bel}(\neg A) - \mathrm{Bel}(\varnothing) = \mathrm{Bel}(A) + \mathrm{Bel}(\neg A)$$

So with belief functions, you only need have: $\mathrm{Bel}(A) + \mathrm{Bel}(\neg A) \leq 1$.

But with probabilities, you must have: $\Pr(A) + \Pr(\neg A) = 1$.

I can believe in A a little bit,
without having to believe in ¬A a lot.

Mass Functions

Belief functions assign numbers to events based on the accumulated "mass of evidence" of their subsets.

$$Bel(U) = \sum_{X \subseteq U} mass(X)$$

One *might have thought* it would be fine just to assign masses to individual worlds (outcomes) and accumulate them like this:

$$Bel(U) = \sum_{x \in U} mass(x)$$

But Shafer's claim is that **evidence is provided at the level of specific *subsets* of W, not at *elements* of W**.
E.g. a sensor reading can provide nonzero evidence for {x,y} while providing zero evidence for {x}, {y}, and ∅.

Mass $\leftrightarrow$ Belief

The form of a sum over all subsets of a finite set

$$B(U) = \sum_{X \subseteq U} m(X)$$

has some nice mathematical properties.
It can actually be inverted this way:

$$m(U) = \sum_{X \subseteq U} (-1)^{|U|-|X|} B(X)$$

This is called a **Möbius transform.**
Evidence mass is the Möbius transform of belief.
Belief is the "inverse Möbius transform" of evidence mass.

The Möbius Transform Defined

Let W be a finite set. Consider an arbitrary function B: $2^W \rightarrow [0,1]$.

Its **Möbius transform** $m = \mathbf{M}(B)$ is another function m: $2^W \rightarrow [0,1]$ defined by:

"differentiating"

$$m(U) = \sum_{X \subseteq U} (-1)^{|U|-|X|} B(X)$$

Let W be a finite set. Consider an arbitrary function m: $2^W \rightarrow [0,1]$.

Its **inverse Möbius transform** $B = \mathbf{M}^{-1}(m)$ is another function B : $2^W \rightarrow [0,1]$ defined by:

"integrating"

$$B(U) = \sum_{X \subseteq U} m(X)$$

The Möbius Transform Calculated

$$m(U) = \sum_{X \subseteq U} (-1)^{|U|-|X|} B(X)$$

This formula is simpler than it looks at first glance.

$m(U) =$ sum of B applied to each subset of size n
$-$ sum of B applied to each subset of size $n-1$
$+$ sum of B applied to each subset of size $n-2$
$-$ sum of B applied to each subset of size $n-3$
$+$ *etc*....
down to $B(\emptyset)$.

The Möbius Transform: Example

Let W = { x, y }.
Let B(W)=1   B({x})=0.6 B({y})=0.1 B(∅)=0
Then:

$$m(W) = 1 - 0.6 - 0.1 + 0 = 0.3$$
$$m(\{x\}) = 0.6 - 0 = 0.6$$
$$m(\{y\}) = 0.1 - 0 = 0.1$$
$$m(\emptyset) = 0$$

Inverting this:

$$B(\emptyset) = m(\emptyset) = 0$$
$$B(\{x\}) = m(\emptyset) + m(\{\mathrm{x}\}) = 0 + 0.6 = 0.6$$
$$B(\{y\}) = m(\emptyset) + m(\{\mathrm{y}\}) = 0 + 0.1 = 0.1$$
$$B(W) = m(\emptyset) + m(\{y\}) + m(\{y\}) + m(W) = 0 + 0.1 + 0.6 + 0.3 = 1$$

What *are* Möbius Transforms of Belief Functions?

$m(U)$ seems to be "the amount of belief committed to U that has not already been committed to its subsets [Halpern 2003, p36]."

E.g. W= {hep, cir, gal, pan} (mutually exclusive syndromes).
Consider with a diagnostic test that "was positive 70% of the time when a patient had hep or cir."
Think of how an m function could represent this test:

$$m(\{\text{hep,cir}\}) = 0.7.$$

Ok. How should this constrain $m(\{\text{hep}\})$? Or $m(\{\text{hep,cir,gal}\})$?
Can they both still be 0?
Shafer: Yes.