Testimony as Evidence

Katie Steele

Department of Philosophy, Logic and Scientific Method
London School of Economics (LSE)

FEW
May 2011

Preamble

Ultimately we are concerned with what 'short-cut' methods for updating on testimony are Bayesian-compatible.

Firstly, however: What is testimony? And why the shortcuts?

On Testimony

- another's beliefs on some issue that concerns us *both*
- representable as a probability distribution across a common partition: e.g. $Pr_T(B_1, B_2, B_3) = (0.2, 0.3, 0.5)$
- may have been gleaned in different ways, e.g. inferred from verbal communication, physical traces
- thus some amount of inference not explicitly modelled...
- ...indeed, further assumption of *reflective equilibrium*...

Reflective Equilibrium

- Assume: agents do not have common priors; not necessarily same proposition space
- Lehrer and Wagner (1981): agents have (informally) traded background evidence regarding **B**, as far as possible
- I prefer: principle agent has settled on a key issue where her views differ from her peer; this is the issue for updating.
- If you like: principle agent is updating in response to a portion of her peer's priors.

How to update on testimony?

We seek plausible candidates for the function:

$F_{\mathbf{B}}$: prior, $M_{\mathbf{B}} \rightarrow$ posterior

Where:

- $M_{\mathbf{B}}$ is the testimony matrix
  –may be more than one 'expert peer'
- prior/posterior refers to initial/final probability functions just on **B**

Agenda

1. Why not Bayesian business as usual?
2. Linear averaging for single partition: quick defence
3. The problem of rich event spaces
4. Testimony as EVIDENCE
5. Concluding remarks

Why not standard Bayesian?

Bayesian model:

Testimony, $M_{\mathbf{B}}$ explicitly part of model. This is what is learnt.

$$Pr'(B_j) = Pr(B_j|M_{\mathbf{B}}) = \frac{Pr(M_{\mathbf{B}}|B_j) \times Pr(B_j)}{Pr(M_{\mathbf{B}}|B_j) \times Pr(B_j) + Pr(M_{\mathbf{B}}|\neg B_j) \times Pr(\neg B_j)}$$

So what is the problem?

–well, nothing really
–but the likelihoods in the above expression are somewhat awkward...
–...hence, the appeal of 'shortcut' methods?

Popular alternative:

Posteriors across **B** achieved via LINEAR AVERAGING

No 'indirect' likelihoods; rather, direct pooling.

$$Pr_0'(\mathbf{B}) = w_0 \times Pr_0(\mathbf{B}) + w_1 \times M_{\mathbf{B}}[1] + \ldots + w_n \times M_{\mathbf{B}}[n]$$

Where: The weights $w_0 \ldots w_n$ are understand as 'weights of respect'; they are non-negative and add to one.

Note that defence of LA given in Wagner (1985) entails same weights, regardless of agent's prior and contents of testimony matrix.

Linear averaging for single partition: quick defence

Quick defence of Bayesian compatibility:

Think of a single episode of averaging.

There is *some* Bayesian model that will represent this, i.e. some model that gives

$$Pr(B_j|M_{\mathbf{B}}) = \frac{Pr(M_{\mathbf{B}}|B_j) \times Pr(B_j)}{Pr(M_{\mathbf{B}})} = Pr(B_j) \times w_0 + M_{\mathbf{B}}[1](j) \times w_1 + \ldots$$

N.B. Even if we retain the Wagner defence of LA, the weights can be different for different expert groups (response to Bradley 2007).

The problem of rich event spaces

But linear averaging is in any case incomplete.

- no use as a 'short-cut' method if it does not give us a *full* posterior function on all propositions
- presumably we also want to do multiple updates on orthogonal partitions

Once we complete the method, however, there may not always be a Bayesian representation where the testimony can be explicitly modelled as learning propositions $M_{\mathbf{B}}$, $M_{\mathbf{C}}$, etc.

Can we get around these issues for Linear Averaging?

An initial proposal:

- Supplement averaging with Jeffrey conditioning
- RE. Bayesian compatibility: consider testimony as *outside* the model; not represented in proposition space. It is an *extra*-Bayesian updating rule.

There are more issues to consider, however...

The problem: Can show by example that two testimony updates by the method just outlined do NOT generally commute.

	B	$\neg B$		B	$\neg B$		B	$\neg B$
C	0.1	0.2	$\rightarrow [B_{0.9}, \neg B_{0.1}] \rightarrow$	0.225	1/30	$\rightarrow [C_{0.6}, \neg C_{0.4}] \rightarrow$	0.523	0.077
$\neg C$	0.3	0.4		0.675	2/30		0.364	0.036

However...

	B	$\neg B$		B	$\neg B$		B	$\neg B$
C	0.1	0.2	$\rightarrow [C_{0.6}, \neg C_{0.4}] \rightarrow$	0.2	0.4	$\rightarrow [B_{0.9}, \neg B_{0.1}] \rightarrow$	0.485	0.064
$\neg C$	0.3	0.4		0.17143	0.22857		0.415	0.036

Testimony as Evidence

Is non-commutativity really a *problem*?

Well I take commutativity to be the hall-mark of incremental evidence...

i.e. new evidence *adding* to old evidence, rather than new evidence *overriding* old evidence.

And why would testimonial evidence not be incremental?

Testimony as Evidence

New proposal. Background:

We exploit Wagner's sufficient conditions for commutativity.
cf. Field (1978), Diaconis and Zabell (1982), Jeffrey (1988).

$\longrightarrow$ updates are commutative with respect to their *identical* counterparts.

Here, *identical* means: same set of *Bayes factors*, i.e. same set of ratios, corresponding to all pairs of propositions, of posteriors over priors.

Testimony as Evidence

New proposal:

If testimony, e.g. $M_{\mathbf{B}}$, produces *identical* learning à la Wagner ...
... then we get commutativity.

So let us make that a property of the updating rule.

This constraint requires that we replace linear averaging by something from this class of functions:

$$Pr'_{\mathbf{B}} = \text{normalize } [Pr_{\mathbf{B}} \times f(M_{\mathbf{B},G})]$$

Testimony as Evidence

Consider an example function from this class:

$$Pr'_{\mathbf{B}} = \text{normalize } [Pr_{\mathbf{B}} \times \sum_{i=1}^{n} w_i \times M_{\mathbf{B}}[i]]$$

That is, we average the testimony, then multiply by the agent's prior.

Let us consider an interesting property of this function (besides commutativity)...

Testimony as Evidence

Main point: agent can only *defer* to another agent, no matter what the other agent says, when they have flat priors for the issue in question.

cf. linear averaging

Example

I may have some opinions on an issue in theoretical physics, **B**.

Then I meet Stephen Hawking, who has a particular probability function on **B**.

I cannot *defer* to Hawking, as I was already opinionated on **B**.

Maximal respect by my example updating function gives:

$Pr'_{\mathbf{B}} =$ normalize [$Pr_{\mathbf{B}} \times$ Hawking function]

Concluding Remarks

Objections?? [...to follow !]

In general, we will need to be careful about *what* counts as *identical* testimony.

Can we fill in these details while still retaining the benefits of a 'shortcut' updating method??