KBNS AP Seminar, 26 March 2019
On Paul Boghossian's 'Delimiting the boundaries of inference' (*Phil. Issues* 2018)

Minuted discussion points are interlaced *[[[triple-bracketed and italicized]]]*.

§1 Introduction

1. Aims and methods

(i) PB's main aims in the paper are:

> (1) To ask, what features any notion of inference must sustain if inference is to be a 'central epistemological notion', or a notion of an important 'mental or epistemic *kind*' (§3);
> (2) To consider which of the many things people sometimes seem to have in mind when they speak of 'inference' will sustain these features.

(ii) Under (1) PB proposes (p 59):

> B. (Basing) When you infer from P to C, you *establish* the premises P as *your reason* for believing the conclusion C, you establish them as the *basis* on which you believe B.
> Q. (Quality) Given this basing fact, your belief can be assessed as resting on good or bad reasons.
> R. (Responsibility) And given this fact, too, you can be held *responsible* for having reasoned well or badly. The quality of your reasoning will enter into an assessment of your rationality.

(iii) PB's approach to (2) is:

> (a) to set out a 'paradigm example' of 'fully explicit reasoning' (FER, p. 56) and to abstract from it a generalized model of what it involves, the 'intellectualized model of fully explicit reasoning' (IMFER, p 57);
> (b) to categorize various conceptions of, or proposals about, what constitutes inference according to which of the features of this model they retain and which they surrender (pp 57-9);
> (c) to argue that, while some aspects of the IMFER model can be surrendered (or qualified), the most liberal or inclusive conceptions of inference in play in recent discussions will not sustain its essential features as an 'epistemological kind' (p 61 for the 'easy' cases, and pp 63ff for the 'harder' cases).

(iv) Obviously this approach – of starting with a clear and central case and seeing what we can subtract from it without losing its crucial epistemological features – is viable *only if* the clear and central case is itself coherent. So PB includes (pp 61-3) a short section responding to criticisms of his way of thinking even about this central case.

Other initial clarifications

(i) The topic is not:

> 'argument' in the logic-book sense, in which an 'argument' is 'just a set [or sequence] of propositions' (p 55),

but:

> inference as 'a movement of thought' (p 55) that results in [or is directed towards?] the formation or adjustment of ppnal attitudes.

(ii) The discussion is meant to be general, in including both 'theoretical inference' – leading to the formation/adjustment of beliefs – and 'practical inference' – leading to the formation/adjustment of intentions. [But PB then attends mostly to theoretical inference (p 56).]

(iii) PB speaks of the 'agential conception' to emphasize that this 'movement' is one we *make* rather than one that simply *happens*:

> inferring has to be something you *do*, and not just something that happens to you. It has to be a mental *action* of yours… (p 60)

[[Comment:
In these notes I'll (probably) just assume that this is in *some* sense right. But obviously 'something we do' is not very helpful on its own. (Tripping over uneven pavements, and landing heavily if dropped from a height, are things people do.)
Maybe the most important issues in PB's paper surround his attempts to make this notion clearer and more discriminating. Questions to have in mind about that from the beginning might include:

> (i) whether the importance attached to *Responsibility* in the way this paper attempts this, so as to establish inference as a 'mental or epistemological kind', fits with his (apparent?) aim in previous papers to establish inference as a 'psychological' kind. (It would be at least contentious to assume that whether someone is responsible for φ-ing, or whether her performance in φ-ing is evaluable by reference to certain standards, is settled entirely by 'what goes on in her mind' when she φ-s.)
>
> (ii) PB seems (?) to assign an important role in clarifying the relevant understanding of 'something we do' to a (Humean?) 'means-end' model of practical rationality:
>
> > The agent has an aim; she has a view about a way of accomplishing that aim; and she performs an action as a result of that combination. (pp 61-2)
>
> Is that the only model, or the right model, and if so how exactly is it to be applied?]]

[[[There was some initial discussion of point (ii) above, and agreement that:
(i) the conception of inference as an action informed by (or directed towards) some end, so that both it and the agent are evaluable by general canons of a means-end model of practical rationality, did appear to have central role in PB's account of it;

(ii) at this stage in the discussion, just what *end inference is reckoned to serve was unclear, and so further discussion might be better left until more of the account was in view.]]]*

§3a The epistemologically crucial features (p 60)

1. Basing

> B. (Basing) When you infer from P to C, you *establish* the premises P as *your reason* for believing the conclusion C, you establish them as the *basis* on which you believe B.

[[Comments / questions:

(a) in previous seminars we've asked whether 'basing' is meant to indicate something 'dynamic' or 'static'. Here it's both:

> 'dynamic', in that it is an occurrence that *establishes* something to be the case;
> 'static', in that what is thus established then just *is* the case.

[[[As it has figured in some other recent seminar discussions the 'static' relation of basing appears mysterious, or at least complicated. In those contexts it seemed at least to include:

> *(i) causal dependence, as when one state or process causally depends on another (e.g. when a weight's sitting in one pan of a balance-scale sustains the position of the other pan (state), or when the brain's pumping oxygenated blood sustains brain activity (process);*

but to exceed that in that

> *(ii) the causal relation is of the 'reason-giving' kind.*

In PB's discussion it appears that (ii), the 'reason-giving' character of the relation, is secured by other aspects of the account:

> *(a), that the relation is 'established' by an inference; and*
> *(b) that this involves the 'taking condition' – that the thinker appreciate (or seem to) that P supports C.*

If the theoretical weight of (ii) can in this way be taken by other parts of the account, perhaps 'basing' as a static-relation can be more straightforwardly understood as involving (i) alone.]]]

(b) 'your reason':

Nothing you do will establish P to be a reason *to believe* C if it is not.

What inference (is said to) do(es) is 'establish P as *your* reason *for believing* C'.

So reasons *for believing* can be good or bad, so count or not count as reasons *to believe*.

But there are limits: one can't, by 'inferring' tomorrow's weather forecast (C) from some arcane fact of Roman history (P), 'establish' that P as one's reason for believing that C. Where do these limits come from?

[[[Intuitively, the limit is the (hazy) limit on what kinds of error are intelligible: mistakes are possible, but for something to count as a mistake *it must at least be intelligible how a thinker might regard it as not mistaken.*

In the case in hand: no one could intelligibly *be thought of as 'inferring' the 'conclusion' C, that it will rain tomorrow, from the 'premise' P, that Nero was given to fiddling at inopportune*

moments, and thus as believing C for the reason that *P, given that P is* so patently not *a reason to believe C.*
(Alternatively, and given that any inference from P to C can be recast as a modus ponens with P→C as premise, the limit appears as a limit on which such conditional premises the thinker can intelligibly be taken to believe.)

It was observed that, if any such rough account of the 'limits' is right, it suggests:
(i) that an account of good *inference, or inferring* well*, needs to be given a central place in an account of what inferring (good or bad) amounts to; that one should not aim* first *to explain inferring (good or bad) and* then *explain the good cases as those that satisfy an extra condition.*
(ii) there is at least a tension *between*

> *(a) the aim to explain 'inferring' as a psychological kind – one delimited by reference to the natural causal transitions and dependencies it involves,*

and

> *(b) recognizing that there are broadly normative constraints on whether something can be an instance of that kind.]]]*

(c) 'establishing':
It's a little odd to say that *by* inferring C from P one '*makes* it true that your belief has an inferential reason…' (p 60): if believing P didn't *already* give me a reason *for believing* C, would I make the inference?
(Maybe the idea is that a *reason for* φ-ing only gets to be a *reason why* I φ-d if I do φ; but that seems too trivial an understanding of 'established'.
Maybe instead the idea is that there are lots of reasons for believing C, maybe lots of reasons I 'have', and making the inference selects the ones on which my belief that C will henceforth be (said to be?) 'based'. If so, I'm sceptical: I've no idea how or when or why (e.g.) I first came to believe that Frege thought that incomplete expressions have reference – but I *could* now *give* reasons for believing that.)]]

[[[Several observations were entered here:
(i) Against the 'sceptical' thought:
it need not be assumed that the 'static' relation of basing established by inference is very long-lasting.
(ii) In keeping with the sceptical thought:
if we accept that the reasons someone could *or* would *now* give *for believing something are more relevant than its history to whether and how this belief is inferentially grounded, this shows at least*
(a) that the epistemologically *relevant 'static' relation of basing is one that* can *obtain independently of being 'established' by an inference once actually performed;*
and therefore
(b) the primary explanandum *should be this* general*, history-independent relation.]]]*

2. Quality

> Q. (Quality) Given this basing fact, your belief can be assessed as resting on good or bad reasons [or: '…is assessable as having good or bad inferential reasons supporting it'].

[[Comments / questions:

(a) 'resting' or (conversely) 'supporting' is intended to be the 'static' relation between attitudes – between one's believing that P and one's believing that C – that is 'established' by inferring. ('Sustaining' might have been a better word.)

(b) Supporting/sustaining reasons are good / bad according to relations between the contents believed:

> …for a reason P for believing C to be good, C must, in some broad sense of 'follow', follow from P. (p 60)

(In *any* broad sense of 'follow', following-from is a relation between the contents of the beliefs, rather than the attitudes of believing.)

(c) Why 'follow from'?
It's stretching a point (at least) to suggest that C (= that the streets are wet) in any sense 'follows from' P (= that it rained recently). But P is, in most circumstances, a reason *to believe* that C.
[It seems (?) that PB deliberately avoids explaining good/bad reasons *for believing* by reference to reasons *to believe*. Why?]

[[[It was observed that this is at least in keeping with PB's aim first to explain inferring (good or bad) as a kind of mental action, the conception of which is not dependent on a prior notion of good inference.
It was not thought that it offered any further, independent reason to question that aim.]]]

(d) 'Following-from' is not sufficient:

> … but that is not enough. FLT follows from the Peano axioms, but it wouldn't do to infer the one from the other… [Why not?] …the conclusion should not be at too far a distance from the premise… [Which means?] …The only good reason I can think of is that the step from premise to conclusion be such that the thinker have some *appreciation* that the conclusion does indeed follow from the premises. (p 60)
>
> (i) It seems that this 'only good reason' is not, in fact, any good.
> As it happens, I *do* (courtesy of Wiles) have 'some appreciation' (in fact, I know) that FLT (semantically) follows from (2nd-order) PA.
> But I don't think this knowledge would, in PB's view, transform my inferring from one to the other into something less bizarre, or something that '*would* do'.

[[[There was extended discussion of this complaint. Point made included:

(i) PB frames the 'taking condition' by emphasizing 'appreciation': 'the thinker must have some appreciation *that…'. This suggests a richer notion than is assumed in the complaint.*
But against that:
(ii) The word occurs, in framing the 'taking condition', in the construction, 'appreciate (or have some appreciation) that *P'; and in that construction its richer connotations are effectively silenced. Appreciating that P is just knowing that P - or, when used non-factively, as PB must* intend *to use it (since some bad inferences will be bad precisely because not-P), it is just believing that P.*

> *(E.g. in debate, 'I suspect you fail to appreciate that so-and-so' is merely a polite way of saying 'You're arguing as you do because you don't know…'.)*

But even if that point is conceded:
(iii) The suggestion of a richer relation is worth exploring, as a modification *to PB's account if not as a reading of it.*
A natural suggestion is that:

> *the thinker should not just think (or believe, or intuit)* that *C follows from P, but* understand how or why *C follows from P.*

This is in keeping with another natural thought about inference:

> *that it should be* the premises themselves*, and not any incidental or external knowledge* about *them, that provide the inferrer with her reason for believing the conclusion.*

(If any such external or incidental knowledge is sustaining the inference, then it should properly be included as an additional premise. In the example given, the thinker would not *then be inferring FLT from PA, but instead from PA together with the further premise that PA entails FLT. And in that case, the inference would be straightforwardly acceptable.)*
(iv) Although this suggestion is worth pursuing, it would be a significant *change to PB's conception:*

> *(a) 'Understanding how or why' is non-negotiably 'factive': one can't understand how or why P unless P.*
>
> *(b) In contrast with 'know' and (the unloaded understanding of) 'appreciate', it is at least hard, and perhaps impossible, to find a non-factive version of it that could figure in an account intended to embrace both good and bad inferring.*
>
> > *(E.g., could it really be a condition on inferring, as an* epistemologically relevant *notion, that someone merely 'have the impression of understanding…'?)*
>
> *(c) So – tentatively – to the extent that this suggested modification is worth pursuing it reinforces the previous reasons for thinking that* good *inference should be the primary explanandum.*

Independently of the above chain of points, there was also discussion of whether the testimonial *status of the supporting knowledge was a relevant factor, either in the oddity of the 'inference' PB mentions, or in the (in)effectiveness of the objection to his view of it.*
(i) A point close to the surface in this part of the discussion, though not distinctly brought out, is that the supporting knowledge is not purely *(or* brutely*) by testimony.*

What the objector knows 'by testimony' (or rather, because it is now just part of general knowledge) is just *that Wiles proved FLT.*
From this 'testimonial' given, an inference is involved, running roughly:

> *Wiles proved FLT → FLT is true; FLT is an arithmetical statement → so an arithmetical truth → so true in all arithmetical structures; PA^2 characterizes those structures; so PA^2 entails FLT.*

(ii) So the case envisaged in the objection is one in which someone 'infers FLT' on the basis of something (that PA^2 entails FLT) that itself is held on the basis that FLT is true.
And it is (to say that least) a little odd to 'infer' a conclusion on the basis that the conclusion is true.
(iii) It is thus worth exploring:

> *whether there are other examples of the intuitive point PB has in mind (a point about there not being 'too great a distance' – or perhaps better, there not being a* rationally inscrutable *distance – between premises and conclusion) that wouldn't raise this complication;*

and if so:

> *whether a parallel complaint could be made against his account of them.]]]*

(ii) I *suspect* that an image of *proving*, rather than of *inferring*, has a big (unacknowledged?) influence in PB's discussion. Lots of actual inferring seems to rest on 'background knowledge' of indefinitely complicated kinds (e.g. applying some fairly-direct consequence of Choice – as originally so-named – in a domain one independently knows to satisfy well-ordering).

[[[Some agreed that the previous discussion – and in particular, the way it naturally introduced assumptions about 'rational scrutability' and 'understanding why', assumptions that have their first home in thought about demonstration *or* proof *– might serve to reinforce this suspicion.]]]*

(iii) It's notable that in this initial discussion this 'only good' [but in fact bad] reason is the *only* ground presented for – and is presented as sufficient on its own to ground – the 'taking condition':

> Inferences require that the thinker *take* her premises to support her conclusion. (p 60)

(iv) PB is concerned at this point to avoid super-tasks:

> If (a) every inference required a taking, and (b) every taking required a prior inference, then no inferring could happen.
> (We are to keep (a) and deny (b): some taking is not a result of inference.)

This is distinct from any worry that *applying* a 'taking' would itself be an inference (p 62).]]

<u>3. Responsibility</u>

> R. (Responsibility) And given this [basing] fact, too, you can be held *responsible* for having reasoned well or badly. The quality of your reasoning will enter into an assessment of your rationality. (p 59)
> [So, w]hat's needed to make it apt to hold you responsible for the quality of your reasoning? (p 60)

[[Comments/questions:

(a) I confess, I'm simply not sure how much (or how) PB thinks this last crucial feature adds to *Basing* and *Quality*.

(i) On the one hand, it is 'not redundant':
- PB points out [rightly] that things people 'do', or their ways of 'doing them', can be assessed as good/bad without that assessment attaching to them as agents (e.g. you're not bad for having bad hearing – but (probably?) are bad for being a bad listener).
- So this condition is 'not redundant on the previous one' (p 60).

(ii) On the other hand:
- the only the only additional factor mentioned here, 'control',

> It has to be a mental *action* of yours, something you have control over, and which you could have done differently, had you thought it desirable to do so (p 60),

seems obviously not to the point (bad listeners can't just choose once in a while to be good listeners); and
- at various later points it seems (?) that it is the 'taking' condition that brings in evaluation of the *agent* rather than only of what it done,

> (e.g. at p 62 PB talks of 'how that taking state could rationally *control* the formation of the conclusion');

and – in PB's view – the 'taking condition' *is* supposed to be implied *already* by *Quality*.

(b) Maybe this doesn't matter much: we *could* agree that the *Responsibility* condition is importantly correct, whether it is an additional condition or not.]]

[[[Provisionally, at least, it seemed reasonable to accept and explore the Responsibility condition, without worrying too much about its relation to the previous conditions.]]]

§2 Candidate conceptions of inference

The job of this section is to try to put into some order a 'fairly comprehensive list of all the types of phenomena that we have been tempted' to call 'inference' (p 59).

The ordering is by reference to a 'paradigm example' of 'fully explicit reasoning':

(1) You consider explicitly some proposition that you believe, for example p.
(2) (Meta Q) You wonder, in the context of some particular inquiry, what other relevant proposition you have reason to believe on this basis?
It then strikes you that q follows from p. Hence,
(Taking) You take it that q follows from p
[Optionally, you might ask yourself:
(3) Is q plausible? Is it less plausible than the negation of p?
(4) You conclude that q is not less plausible than not-p.]
(5) So, you judge q. (You add q to your stock of beliefs.)

[[Comment: The 'optional' bit seems important. Inference, as portrayed here, seems to be aimed at discovering one's commitments in believing that P. Assuming that those commitments are 'wide-scope', then one might say either:

(a) 'inference' has done *its* distinctive job once the 'taking' point is reached, mid-way through (2), and what happens then (3-5) is the balancing of (apparent) *reasons*, not more *reasoning*; or

(b) if we take 'inference' to include more – to describe a transition that reaches all the way to one's new attitude-set – then the 'optional' bit is an essential part of it.

Either way, 'optional' doesn't look to be an optional part (both *a part* and *optional*). Nonetheless, PB omits it to arrive at…]]

[[[There was (inconclusive) discussion of this point.
(i) It was agreed that the choice between (a) and (b) – i.e. between

(a) a conception of inference as consisting only in discovering one's commitments, so that the 'optional part' of determining the best way of meeting those commitments isn't a proper part of it;

and

(b) a conception of inference as including also 'following through' on those commitments, and reshaping one's attitudes to accord with them,

could be regarded as terminological.
(ii) But also agreed, that the choice made here would affect the (still unresolved) question of the 'aim' or 'end' of inference, considered as a rational action.]]]

<u>Inference 2.0, or 'IMFER' – the 'Intellectualized Model of Fully Explicit Reasoning'</u>
In which a thinker:

a. Explicitly judges the premises of the inference;
b. Explicitly wonders, in the context of some particular inquiry, what else she has reason to believe, what other proposition the premises support.
c. Explicitly *takes* the premises to *support* the conclusion.

d. Knows the properties of the premises in virtue of which they support the conclusion (knows the *epistemic principle* that validates moving from the premises to the conclusion).
e. *Believes* the conclusion *because* she believes the premises.
f. *and* believes the conclusion *because* she *takes* the conclusion to be supported by the premises.

[[Not sure I understand this:
- Regarding (b): it looks like a case of 'wondering wh…'; as such it seems it would be brought to its end by 'knowing (or believing?) that…'; that it does reach *some* end is presupposed at (c); but whatever end is reached appears to be *not yet* a case of 'knowing (or believing) that…'.
- Regarding (d): presumably one can know that in virtue of which A only if A; so this condition can hold only if the premises *do* support the conclusion, i.e. it can only be part of a model of *good* inference; but PB intends to model inferring, good *and* bad; so this clause needs to be reformulated – and I'm not sure how.]]

Inference 1.75
Involves (a)-(c) and (e-f); but not (d).

> …he just knows that it seems to him that the premises support the conclusion. (p 57)

[[Is this mis-stated? (It's not obvious why knowing something *about himself* should suddenly be relevant.) What's intended (I think) is: it seems to him that the premises support the conclusion, but he couldn't in general terms state why they do – see p 63.]]

Inference 1.5: Quick and effortless inference
PB summarizes: 'in this case (a) and (e) are both true; but (b), (d) and (f) all seem false' (p. 57).
What about (c), the 'taking' clause?
At this point we're told only that 'it would seem [that there is no] taking state' – so seemingly not-(c); nor therefore does the thinker believe the conclusion because (c) holds – hence, not-(f).
But an adult thinker 'would certainly be capable of thinking' the 'taking thought'; whether that capacity is enough (in PB's estimate) to restore (f) isn't (yet) said.
(Later (p 64) it will turn out that a taking state is 'present *tacitly*'.)

Inference 1.25: Inference in children
On (a)-(f), this scores the same as Inference 1.5.
It subtracts even the capacity to frame the 'taking thought'.

Inference in non-human animals
Not scored against (a)-(f) (but presumably would pass a version of (a) and (e) with 'explicitly' removed?).
[[So far as I can see, the paper doesn't take this category at all seriously.

The only kind of example considered (taken from Kornblith, discussed at p 58) is conditioning, or 'learning' as behaviourists conceived it.
That seems just to be irrelevant to the topic in hand: it has to do with changes over time in how a creature responds to the same 'premise' (or stimulus), not with the nature of the response on any one occasion.
There are certainly more relevant and more plausible examples.]]

'Unconscious and sub-personal inferences', e.g. of the visual system, and 'inferences and computations in computers' etc.
Again, not scored against (a)-(f).
- Re sub-personal (e.g. visual system) inferences: 'we don't hold the person responsible for them' (p 61).

> [[True enough. But since 'sub-personal' means that we don't attribute the inference to the person *at all*, it is hardly relevant. The more interesting questions, whether we hold the *system* responsible, and if not why not, aren't even raised.]]

- Re 'machine inferences': again, 'we don't hold machines responsible'.

> [[John Cleese knows otherwise. But if *mostly* we don't, the (a)-(f) criteria aren't used to say *why* not.]]

[[Comment:
- While the 'inferences' of sub-systems and machines are of little interest in themselves, it seems (to me) to be a major flaw in the design of the paper that PB in these cases departs from (what I've taken to be) his own recommended method, and just doesn't consider how they stand vis-à-vis (a)-(f).
- The overall stance of the paper (I think) is that it is the *taking* condition (c), (and perhaps also the 'aim-setting' meta-wondering at (b) – though I don't understand that), that primarily anchors *Responsibility*: *it* is the condition that makes it 'apt' to evaluate the *agent*, and not just the performance.
-*Why* that should be so I'm not sure: Why are *we* any more accountable for our 'meta'-stances than for our 'first-order' stances?
- But *if* it *is* so then it seems that, where *Responsibility* is obviously not met, its anchor, taking, ought to be (obviously?) missing too.
- PB doesn't test whether this is so, and so far as I can see it is not so: when one adopts the 'intentional stance' towards machines and sub-systems, it comes no less naturally to ascribe them 'taking' stances than 'first-order' attitudes like beliefs and wants (e.g., 'the visual system takes the relative size of a repeated pattern to indicate distance' – hence the 'visual cliff').
- More than that, it seems that the two come together: there'd just be *no point* attributing to the system 'first-order' attitudes to the 'input' information unless one was also ready to attribute to it dispositions to deliver 'output' information on that basis – ways of 'taking' the input 'premises' to indicate or support output 'conclusions'.

- I'm not sure that, in *this* respect, machines are different from us: 'taking' stances would then not be a magical *addition* to believing – one that can anchor agent-evaluations that would be groundless without them – but part of having beliefs at all.]]

§6 Attempting to do without aims and takings (either 'explicit' or 'tacit')

1. Two such attempts *are* considered:

(a) Hilary Kornblith: 'inferences are not mere causal transitions; rather they are *content-sensitive* transitions' (p 64).

This is quickly dismissed because it over-generates: it would let back in sub-personal systems and machines.

> [[Allegedly, anyway. But some would say that computers, and probably bits-of-brain, are sensitive to the grammatical shape of the 'representations' they compute, not their meaning.]]

(b) Mark Richard: responsibility is important, but it's wrong to think that 'one can only be held responsible for things one *does directly*' (p 64).

[[There may well be an interesting idea here. Re the example on p 65:

> - An electrician is skilled, as I am not, in knowing where to look for a fault when a circuit does not work as it should.
> - He *might* run through, and even say out loud to an apprentice, some steps. (But not *all* the steps – whatever that would mean. One can imagine:
>
> > 'A shower-circuit is carrying about 40A; a loose connection makes for resistance, and will more easily burn out the heavier the current running through it; the weakest screw-terminals are in the cheapest fitting, the ceiling-mounted pull-switch; so look in the pull-switch.'
>
> But *some* important 'steps' – e.g. that 'watts-over-volts = amps' – would surely *always* remain as unspoken background to the reasoning.)
> Probably, though, he won't run through steps at all, but will just look to where the fault (probably) is.
> - I, by contrast, might 'think it through', and find the fault, eventually, but only after looking in two or three wrong places.
> - The electrician rightly gets credit not due to me.
> (This electrician doesn't have observational knowledge of the fault; and nor is it a case of Alan's ('deer-tracks') knowledge-from-indicators; but there are similarities: witness a good mechanic's diagnosis of an engine fault just by listening to it – which falls between, more similar to both than either is to the other.
> Suggested moral: much of what PB is constrained to think of as 'inference' is better thought of as competence – a broad, flexible, but nonetheless domain-specific competence – in handling information.)

But PB doesn't treat it at all sympathetically:

- In his portrayal, the electrician just 'finds himself' thinking, 'It's probably the pull-switch' – an imposition from he knows not where.
- Any 'credit' due to him would then have to derive from some retrospective reconstruction of the steps he didn't run through, and not (as seems more credible) from his not needing to run through them.]]

2. Probably more relevant are the (kinds of) attempts *not* considered: attempts to model (moderately) explicit reasoning without a 'taking condition'.

§§ 5&7 Tacit takings

1. For reasons good or bad, PB is persuaded that the taking condition (c) – and again, the aim-setting clause (b) – is essential to inferences that will support *Responsibility*.

> [But in] inferences 1.5 and lower…there seems to be no explicit aim to guide the mental transitions, and no explicit taking state to capture the thinker's view of how the aim is to be achieved. In the absence of such goal-directedness, how are we to explain why 1.5 style transitions count as inferences in our sense…? The answer is that I don't believe we can. (p 63)

2. So, these necessary states have to be present *somehow*:

> …although such states are not present explicitly, they are present *tacitly*…
> [Inference, like much else] involve[s] guidance by tacit states. (p 64)

3. What does tacit guidance amount to?

> The idea of a tacit state is the idea of a mental state that seems not to be an occurrent state of consciousness yet is still present and plays a role in guiding behaviour. (p 66)

[[So far, clear enough. This description fits belief in its usual ('dispositional') sense: I don't have to 'tell myself' that the Co-op sells tobacco; I know it does, and walk there when I want some.]]

4. Can we say more? What PB offers is in the form of an example:

> When we first learn to use an iPhone, the rules for operating the machine are learned explicitly, and followed explicitly. However, after a certain period of use, the procedure…may become automatic, unlabored and unreflective… It's not as if your grasp of the rules has disappeared; if it had, you wouldn't be able to operate the phone… However, the grasp is no longer an explicit item of consciousness, but rather serves to guide the behaviour tacitly… The same goes for playing the piano, using a language, […etc.] (pp 66-7)

[[Comments:

(a) The last sentence is remarkable. *No-one* uses a language by having once learned explicitly the rules for using it; no-one knows what those rules are. Even more obviously, being able to play the piano is not a matter of having 'internalized' an explicit set of instructions – though talking in piano-teaching does help things along.

(b) Incidentally, PB speaks as someone of a certain age – someone who thinks smart-phones etc *do* come with the kind of instructions that I, of a similar age, believe *ought* to accompany them. (In fact the instructions typically read: 'Switch it on and prod around to see what happens'.)
(c) Still, there are cases – Delia's 'cookery' books telling you how to boil an egg – that more nearly fit his description. Will such examples do?
(d) Surely not. In any given inference the 'taking' condition is that the thinker take *these* premises to support *this* conclusion. That is probably something the thinker has never before thought or had occasion to think. It's not something she has learned explicitly and got used to.
(e) *If* people learned to think in logic classes, PB's description might give an account of their *general* inferential dispositions. They don't.]]