THE VAGARIES OF ACTIONS AND THE VERITIES OF MEANINGS

(Revised draft, May 2014; comments invited)

Paulo Faria

This paper is an exercise in the epistemology of reasoning. Its topic is a *prima facie* improbable one: the role of chance in recognition of deductive validity. Let me begin by explaining what I have in mind.

It is one thing to know that a pattern of inference, say *modus ponens*, is truth-preserving; it is quite another to recognize a particular inference as valid. When I speak here of 'recognition of deductive validity' what I have in mind is the latter: it is knowledge of the logical properties of particular inferences that is my concern here. I don't mean the kind of reflective knowledge that a logically sophisticated reasoner would express by making use of such terms of art as 'valid inference' or 'logical consequence'. I mean the ability to correctly take a set of beliefs as providing sufficient reason for a further belief, or a change of mind.

I borrow my title from David Kaplan, who writes in the 'Afterthoughts' to *Demonstratives:* 'Logic and semantics are concerned not with the vagaries of actions but with the verities of meanings.' (KAPLAN 1989: 584-5) In the section of that work to which this sentence belongs, Kaplan is explaining why his investigation of the semantics of demonstratives and other indexicals concentrates on *occurrences of expressions in a context* rather than on the pragmatic notion of an *utterance of an expression by an agent in a context.* He writes: 'Utterances take time, and are produced one at a time: this will not do for the analysis of validity. By the time an agent finished uttering a very, very long premise and began uttering the conclusion, the premise may have gone false. Thus, even the most trivial inference, *P* therefore *P,* may appear invalid.' (KAPLAN 1989: 584) For the purposes of logical regimentation, such accidents of context-dependency will be

circumvented by referencing, say, all occurrences of 'now' to a single instant, all occurrences of 'here' to a single place, and so on. Likewise, all occurrences of the same demonstrative, say 'this', will be referenced to a single object by assuming (what Kaplan calls) a single *directing intention*. We are, that is, invited to think of two different occurrences of 'this' as driven by a single intention aiming twice at the same object, an intention we would presumably express through such behavior as pointing, staring fixedly at the object, and so son. That should do the trick as far as logic is concerned. 'But', asks Kaplan, 'does it leave our logic vulnerable to a charge of misrepresentation? What is it that we hope to learn from such a logic? (...) To assume that one intention can drive two occurrences of a demonstrative seems to me more falsification than idealization.' (KAPLAN 1989: 590)

The distinction between the vagaries of actions and the verities of meanings is supposed to help with these perplexities; yet Kaplan remarks, at the close of his discussion: 'There is something I'm not understanding here, and it may be something very fundamental about the subject matter of logic.' (*Ibid.*)

I will eventually submit a conjecture as to what it is that Kaplan may not be understanding here; also, whether it is indeed something fundamental about the subject matter of logic. But first things first.

My main question is: can we know by reflection alone whether an inference is valid? And then, however the main question is answered, can rightness of reasoning be, if on occasion, a matter of chance — of one's being, as it were, in the right place at the right time? Can it be that whether we succeed in inferring correctly depends on circumstances beyond our ken? And if so, and then turning full circle, how can we know by reflection alone, if indeed we do, whether an inference is valid?

I will answer affirmatively, thought not unqualifiedly, to each of the first three questions. That is, I will argue that to a *very large* extent

it is by reflection alone that we know whether an inference is valid. But I will also argue that rightness of reasoning can be a matter of chance; that whether we succeed in inferring correctly may well depend on circumstances evading our knowledge; *and* that this is no hindrance to our being, more often than not, able to know by reflection alone whether an inference is valid.

That will leave us with a picture on which the contingencies of external individuation of thought contents (no matter how widespread the phenomenon should turn out to be; no matter, in other words, whether anti-individualism is true) will *not* put in jeopardy reflection as a mode of access to logical form. But it is also part of the picture that reflection may not be enough: that we may sometimes have to count on what John McDowell calls a courtesy by the world.

The main question (that of what Paul Boghossian described as 'the apriority of our logical abilities' (1992), has been usefully - if, to my mind, unsatisfyingly - discussed in the framework of the still-ongoing debate about the compatibility of anti-individualism and first-person authority. It is easy to see why: if many, perhaps most, representational mental states and events depend for being the specific states and events that they are on non-representational relations between the individual and a wider environment, then even fully conceptualized contents (those we ascribe in oblique, *de dicto*, clauses) would seem to fail the intuitively plausible requirement of *transparency*, thus defined:

> The content c of a propositional attitude a of S is transparent to S iff for every content c^* S is in a position to tell by reflection alone whether $c=c^*$.

So construed, transparency requires both identity and difference of contents to be accessible to reflection. The external individuation of content would seem to jeopardize that access, paving the way for failure to recognize either identity or difference of content. The predicament of Kripke's Pierre is an instance of the former failure; that of the slow-

switched denizens of Twin Earth, like Boghossian's Peter, an instance of the latter.

The threat to transparency of validity is now manifest. Here is the scheme of a putatively valid reasoning:

1. Fa
2. Ga
3. $\exists x(Fx \wedge Gx)$

Now add subscripts to your non-logical constants, if only to mark that they are different tokens of the same type, hence not *logically* guaranteed to be co-referential (we are supposed to be dealing here with some real-life inference, not with logical idealization):

4. Fa_1
5. Ga_2
6. $\exists x(Fx \wedge Gx)$

And there we are: in *Environment 1,* $\delta('a_1')=\delta('a_2')$, and the inference is valid. Switch to *Environment 2*, where $\delta('a_1')\neq\delta('a_2')$, and the inference is a fallacy of equivocation. And now suppose you either raise in Environment 1 the possibility that you are in Environment 2, and accordingly refrain from drawing the conclusion, or else you go through the steps 1 to 3, taking yourself to be in Environment 1 when, in fact, you are in Environment 2. Either way, your reasoning abilities are crippled.

The second alternative (you take yourself to be in Environment 1 when, in fact, you are in Environment 2) is, of course, the template exploited in the "slow switching" thought experiments, and in a good deal of the extant literature about recognition of validity. The first one (you raise in Environment 1 the *possibility* that you are in Environment 2, and refrain from drawing the conclusion) will resurface nearing the end of my discussion.

Now I want nothing with slow switching and what Boghossian (1994) called the 'semantics of travel'. I think these scenarios of

metaphysical victimization (you go to bed at home and, while sound asleep, are stealthily taken overnight from Earth to Twin Earth without ever realizing it) have had a corrupting influence on the minds of those who have written on knowledge of validity. They have, in particular, made all but invisible the ways in which, in real life as opposed to philosophical fiction, the vagaries of actions are apt to affect the verities of meanings.

Look at the matter this way: given the task of assessing the inferential behavior of someone who performs the inference 1-3 in Environment 2, wholly unawares of being in such an environment, you are faced with a trade-off between what Joseph Camp calls doxastic and inferential charity (CAMP 2002: 38-39). That is, you will have to choose between maximizing the subject's set of true beliefs at the expense of her rationality or else maximizing her rationality at the expense of her set of true beliefs. Other things being equal, faced with someone who performs the inference 1-3 in Environment 2, you go for inferential charity and look for a (possibly tacit) false premise.

Why is that so? Because you are aware of the unavailability, for the subject whose reasoning abilities you are assessing, of the relevant information about her environment. She has just done her best given the available information.

And that's what I find so unsatisfying about slow switching scenarios and the way the debate about 'the apriority of our logical abilities' has been carried on: the switched subject is *stipulated* not to have cognitive access to the shifts in environment which are apt to affect the validity of her inferences. Here is Tyler Burge, and bear with me to have him quoted at some length:

> Suppose that one underwent a series of switches back and forth between actual earth and actual twin earth so that one remained in each situation long enough to acquire concepts and perceptions appropriate to that situation. Suppose occasions where one is definitely thinking one thought, and other occasions where one is definitely thinking its twin. *Suppose also that the*

switches are carried out so that one is not aware that a switch is occurring. The continuity of one's life is not obviously disrupted. So, for example, one goes to sleep one night at home and wakes up in twin home in twin bed – and so on. (Your standard California fantasy.) Now suppose that, after decades of such switches, one is told about them and asked to identify when the switches take place. The idea is that one could not, by making comparisons, pick out the twin periods from the "home periods". (BURGE 1988: 115. *My emphasis, PF*).

As we should expect, slow switching will have a bearing on 'any reasoning that takes place over time, hence any reasoning' (BURGE 1998: 363). For one thing, the subject's ability to assess rightness of inference would seem to be jeopardized by his unawareness that he's been switched.

That was the problem raised by Boghossian in 'Externalism and Inference': anti-individualism 'is inconsistent with the thesis that our thought contents are epistemically transparent to us (...) this is true in a sense that falsifies another important and traditionally held view – that we can detect *a priori* whether our inferences are logically valid or not.' (BOGHOSSIAN 1992: 13)

And, to be sure, that would be very disturbing indeed. After all, the main interest (and, just possibly, the main promise) of anti-individualism lies in its acknowledgement of the impact of exposure to changing contexts on the constitution of thought contents. That's precisely what explains the interest aroused, in the literature about anti-individualism, by cases of context-switch, often illustrated with such elaborate fantasies about space-travelling from Earth to Twin Earth and back, interplanetary abductions and like exercises in science fiction. At the end of the day, such exercises should have been just a device to graphically describe a range of much less extraordinary phenomena which, if anti-individualism is right, take place in a variety of situations prompted by exposure to differences between the contexts in which the rational capacities of a single subject must be exercised.

Anyway, that's the setting of Boghossian's problem about anti-individualism and inference (BOGHOSSIAN 1992). Boghossian's argument has the form of a *reductio ad absurdum*: the truth of anti-individualism implies the possibility of undetectable errors in reasoning, due to unperceived shifts in propositional content. That possibility clashes with the transparency of mental content; therefore, anti-individualism is false.

Suppose that, having enjoyed a happy childhood on Earth, I am someday carried away to Twin Earth. And suppose further, as Burge asks us to, that 'that the switches are carried out so that one is not aware that a switch is occurring. The continuity of one's life is not obviously disrupted.' (BURGE 1988: 115). I just wake up on Twin Earth in twin bed and everything looks exactly as before. (Remember, that was built into the thought experiment). In due time, according to anti-individualist common wisdom, my usage of the term 'water' comes to mean what the linguistic community to which I now belong uses it to mean: namely, *XYZ*. And here I am, inferring from the conjunction of true premises

7. I enjoyed playing in water (=H_2O) when I was a kid *(from memory)*

8. This glass is full of water (=*XYZ*) *(from current perception)*

the false conclusion

9. This glass is full of the same liquid I enjoyed playing in when I was a kid.

To make things worse, the fallacy I fall prey to is not comparable to the usual fallacy of equivocation, in which an ambiguity is neglected, and the reasoner is in principle in a position to detect and rectify, on a wholly *a priori* basis, the flaw in her reasoning. In the slow switching scenarios, there is just nothing the subject can do to prevent or fix the irrationality, short of undertaking an empirical investigation of the environment, and of her own personal history.

Worse still, given the conditions built into the slow switching stories, it is (to put it mildly) unclear what such an 'empirical investigation' could possibly look like. It's not as if there might be traces, like a flight ticket from Earth to Twin Earth in the inside pocket of your jacket, or custom papers attached to your passport, or a message from home in the answering machine. The two worlds were stipulated to be indiscernible, exact duplicates (apart from the single 'external' difference): were it not so, we would have learnt nothing from the fictions.

The perception that there's not much that the victims of slow switching can do to detect the external sources of their possible logical shortcomings – a perception which I find to be widespread (if mostly tacit) in the literature – answers for the surprising willingness, displayed by friends and foes of anti-individualism alike, to devise exculpating moves as a response to those scenarios of logical misfortune.

After all, the *only* difference between the lucky and the unlucky reasoners lies wholly beyond their ken. I suppose Boghossian speaks for most writers in the field when he writes: 'It seems to me that there is an immediately recognizable sense in which there can be no difference *in respect of rationality* between *[the reasoner in Possible World 1 and that in Possible World 2]*. It seems implausible in the extreme to say that they differ in their capacity to *reason*'.' (BOGHOSSIAN 1992: 27).

That is of course an eminently plausible appraisal, given what the differences between the two contrasting worlds are supposed to be. Small wonder, then, that the choices on offer are (with the remarkable exceptions of SORENSEN 1998[1], and WILLIAMSON 2000) a variety of exculpating moves, designed to shield the rationality of the possibly unlucky reasoner against the contingencies of context-shifting. (See,

---

[1] On which see Faria 2009.

e.g. SCHIFFER 1992, BURGE 1998, LUDLOW 2004, SOSA 2005, COLLINS 2008, RECANATI 2012).

The *individualistic* exculpating move will predictably recoil from the broader and unsafer landscape to an inner domain, a.k.a. narrow content, sealed off from the contingencies of external causation, accident, and luck. Here is how Boghossian introduces it: "If, then, it is also true that there is an important sense in which *[the reasoner's]* behavior *makes sense from his point of view*, we would appear to have here an argument for the existence of a level of intentional description which conserves that sense'. (BOGHOSSIAN 1992: 28).[2]

A bit more surprising are the *anti-individualistic* exculpations, paramount among which is the Schiffer-Burge "anaphoric" view of content preservation.[3]

The main idea here is that the reiteration, in an occurrent episode of thinking, of the content of a past thought is made possible by a dependency relation comparable to that which holds between relative pronouns, and other anaphoric expressions, and their antecedents in the linguistic constructions in which they feature. In 'Laura was confident that she would get the prize', the pronoun 'she' designates Laura: its semantic value is determined by the anaphoric antecedent which is the proper name – as the value of a bound variable in first order quantification is determined by the quantifier which is its anaphoric antecedent. Hence, the unhappy reasoning 7-9 would get reinterpreted as:

10.   I enjoyed playing in <u>water</u> when I was a kid

11.   This glass is full of <u>water</u>↑.

12.   This glass is full of the same liquid I enjoyed playing in when I was a kid.

---

[2] The move bears comparison with Kant's forceful "shrinking" of the proper domain of moral assessment to the inner realm where a pure will operates by itself, sealed off alike from the vagaries of causation, contingency and luck. See Bernard Williams' perceptive discussion in WILLIAMS 1976.

[3] See SCHIFFER 1992, BURGE 1998.

Here I resort to '<u>water↑</u>' to mark the anaphoric dependence of 'water', as tokened by the reasoner in premise 11, on its occurrence in premise 10. And what we have as a result is, small wonder, a valid argument with a false premise: as 'water' in 10 denotes H_2O (the thought content here being supplied by preservative memory), premise 11 amounts to the false judgment that the Twin Earth glass is full of H_2O.

A more surprising anti-individualistic path to exculpation is provided by Peter Ludlow's "Orwellian" theory of content preservation. On Ludlow's theory, 'it is not the job of memory to record contents, but rather to provide information about past episodes relative to current environmental conditions.' (LUDLOW 1998: 316).[4] As in Orwell's *Nineteen Eighty-Four*, the past is rewritten from the standpoint and priorities of the present. Hence when I, on Twin Earth, recall having played in water as a kid, the content of my "Orwellian" memory is the false judgment that I played in twater (*XYZ*). The outcome is, again, a valid judgment with a false premise:

13.   I enjoyed playing in water (*XYZ*) when I was a kid. (*False*)
14.   This glass is full of water (*XYZ*). (*True*)
15.   This glass is full of the same liquid I enjoyed playing in when I was a kid. (*False*).

Now there's no question that all these different construals somehow manage to take into account the fact that, as Boghossian says, the reasoner's behavior 'makes sense from his point of view'. What is not so conspicuous is the sheer absence, *built into the very terms of the slow switching thought experiments*, of any other point of view against which the subject might try and assess the soundness of her reasonings. And that is my complaint against the freewheeling use of so many thought experiments in contemporary analytical philosophy: we end up losing some of our grip on what things look like in some real life, down to earth surroundings.

---

[4] The theory is further articulated in LUDLOW 2004.

The shared assumption underlying all the exculpating moves we have considered is explicitly stated by David Sosa: *'Ignorance is insufficient for incoherence*: inferring subjects are in principle in a position to avoid invalidity, no matter what their state of knowledge (indeed, no matter what the truth of their premiss beliefs.).' (Sosa 2005: 219).

As the attentive reader will not have failed to notice, there is a further assumption at work here: namely, that ignorance is always excusable. Which, as I stressed, makes perfect sense in the contrived setting of the slow switching thought experiments. *There* was indeed virtually nothing that the switched subjects could do to prevent the fallacies of equivocation they were prone to – hence the appeal of the exculpating moves we briefly reviewed.

Against that tide, I want to avoid the pressure towards exculpation, and the felt urge to protect the subject's rationality at the expense of her true beliefs. But then the first thing is to get rid of science fiction and follow Wittgenstein's advice, bringing words back home.

What happens, then, if we detach the examination of the main question from that framework? To begin with, we will stay on Earth and take into account the possibility that the information which the subject actually lacks about her environment *is*, after all, available; moreover, that the subject would be apprised of it had she only cared to know.

No 'semantics of travel', then: uncontroversial examples of external individuation of content - beginning, foreseeably enough, with singular thoughts - should suffice for our purposes. If anti-individualism is right, then the opacity of validity will be just more widespread than we need to assume for the purposes at hand.

Suppose then as I'm coming home in the afternoon I notice this beautiful Golden Retriever dog playing around in my neighbor's front

yard. I stop for a while to pet my new acquaintance, who turns out to be very amiable. As I walk home I think 'That's a very friendly dog.'

Now, a couple of days later, same scene – or so it seems. Here's the nice front yard with its blooming bushes and this beautiful golden dog running around. Again I stop, hoping to attract the dog's attention, resorting perhaps to whistling or finger clapping, yet this time to no avail: the dog keeps running nonstop around the yard, heedless of my inviting moves, barking up every other tree. Maybe he's spotted a cat, who knows. I walk home thinking: 'That's a very restless dog'.

Am I now entitled to infer that there is a dog in my neighborhood who is both friendly and restless?

16.   Fa

17.   Ga

18. $\exists x\ (Fx \land Gx)$

Well, suppose my neighbor is a breeder of Golden Retrievers and what I successively spotted on those two occasions were a pair of siblings from the same litter – call them Argos the Friendly and Targos the Restless. As things go, Argos is not excitable at all, while Targos is of a rather unfriendly disposition. Suppose further there are no other dogs in the neighborhood. So my conclusion is just false, and my reasoning is invalid – a plain fallacy of equivocation. Its form is not 16-18 but rather:

19.   Fa

20.   Gb

21. $\exists x\ (Fx \land Gx)$

And the trouble lies in the way my mistaken empirical assumption (that there was one single dog which I encountered twice) impinges on my grasp of the logical form of the inference I performed – specifically in my taking 19-21 to be of the form 16-18.

For it's not as if I inferred validly, except that my inference relied on a tacit (and false) identity premise (namely that 'that dog$_1$' = 'that dog$_2$') – so that my reasoning was really an enthymeme:

22. Fa
23. Gb
24. a=b
25. $\exists$x (Fx$\wedge$Gx)

That will not do, for at least two reasons. First, it's not as if I would be thinking of each dog as, say, 'the dog I saw on Tuesday' and 'the dog I saw on Friday' and then raising and answering an identity question. On both occasions of meeting with the dogs, the demonstrative concept 'This dog' did all the job of referring.

Sure I could any time introduce two separate descriptions, say 'The dog I spotted on Tuesday' and 'The dog I spotted on Friday', and raise an identity problem; but why would I do that when I'm utterly heedless of the possibility that I met with two different dogs?

But there's worse. Suppose, to raise the problem in a wholly general setting, at t_1 I see object a_1 and think 'This is F'. At t_2, I see object a_2 and think 'This is G'. Then I draw the conclusion: 'Something is both F and G'. Am I entitled to that conclusion? Well, of course, provided $a_1=a_2$. But was that a tacit premise, and my inference an enthymeme?

Kaplan considers the possibility with respect to occurrences of pure indexicals. He remarks that 'You stay. Therefore, it is not the case that you do not stay' is not as it stands an instance of the law of Double Negation (as the references of the two free-standing occurrences of 'you' are left indeterminate); he then wonders whether that could be fixed by an effort to have both occurrences of 'you' refer to the same individual - fixing one's attention and trying not to blink in the meantime, as he says. But if so the form of the argument would really be 'You$_1$ stay. Therefore, it is not the case that you$_2$ do not stay', which is not valid.

'Perhaps' he suggests, 'we should give up on Double Negation, and claim that the argument is a valid enthymeme with the implicit premise "You$_1$=you$_2$", the premise we strove to make true by fixing our attention. "All right", said the Tortoise to Achilles, "repeat the argument and this time remember to utter the additional premise".' (Kaplan 1989: 589)

Now in the original Lewis Carroll scenario, Achilles had to face a regress of *inference rules*: the additional premise was at each step the statement, in the form of an ever more complicated conditional, of the rule authorizing the detachment of the conclusion in the preceding step. Here, by contrast, we face a regress of *empirical assumptions*: what is at stake is not which inference rule we are supposed to follow but whether it applies to the case at hand.

Suppose then

26. Fa_1

27. Ga_2

are not enough to for you to infer '$\exists x (Fx \wedge Gx)$'. After all, you need to make sure that 'a_1' and 'a_2' are co-referential. That is, you need the further premise:

28. $a_1 = a_2$

All right, said the Tortoise to Achilles: that's not going to do, either. For now you have to make sure that 'a_1', as it occurs in 26 and 'a_1' as it occurs in 28 are also co-referential; and ditto for 'a_2' as it occurs in 26 and 'a_2' as it occurs in 28. At which point it is manifest that, as Kaplan implies, you are embarked on a vicious regress.[5]

The regress is only stopped if at some point you can just *take for granted* that two tokens of the same type have the same semantic value. But if you are to be entitled to do that at some point into the regress,

---

[5] The full-fledged regress argument is due to John Campbell in a paper antedating by two years Kaplan's 'Afterthoughts' (see CAMPBELL 1987; and, for a brief restatement, CAMPBELL 1994: 75-6).

then you may as well be entitled to do it from the very beginning, so that no regress arises in the first place.

That does not mean, let me hasten to add, that identity of reference will thereby be secured, only that it will be taken for granted (or, as John Campbell says, *traded upon*) rather than taken as a premise, whether tacit or explicit, in the argument. Indeed, the very point of the regress argument lies in drawing the line between tacit premises in enthymematic reasoning and the kind of background empirical assumptions which we are bound to *riskily* take for granted in assessing deductive validity.

An identity statement will be taken as a premise, whether tacit or explicit, only if there is a difference *recognizable from the first-person perspective* (in other words, by reflection alone) between the ways the object is thought of in each of a pair of premises. That's why, should I reason from 'The dog I met on Tuesday is very friendly' and 'The dog I met on Friday is very restless' I would be helping myself to the identity 'The dog I met on Tuesday = The dog I met on Friday' as a premise.

Mark that it's not as if, by contrast, 'This dog', as tokened in the presence of Argos, would be an exercise of the *same* concept as that exercised in a tokening of 'This dog' in the presence of Targos. Having a different extension is indeed a sufficient condition for the two tokens of 'This dog' being exercises of two different demonstrative concepts or modes of presentation. But the difference, most certainly here anyway, ain't in the head. If there are two Fregean senses at play here, they are *de re* senses, their identity partially fixed by the subject's non-representational relation to his environment. (And, again, if anti-individualism is right, then the problem is just more widespread than, for the purposes at hand, I am assuming here.)

That is what makes logical appraisal, if on occasion, a matter of chance - of one's being in the right place at the right time. It's not just what we *aim at* referring to; it's what we *succeed* in so doing; and this is

not wholly up to us. That's why Kaplan's appeal to directing intentions is ultimately frustrating in just the way he suggests it is. Sure I must aim at a particular dog when I think 'This dog is friendly'; but what dog (if any) I succeed in meaning is a matter of how things are in my environment, and that may well lie beyond my ken.

But now it is high time that we pause to ask: is that really any news? What, indeed, is the fuss all about? Isn't the effect of detaching the question of the 'apriority of our logical abilities' from its framework in the debate about anti-individualism and self-knowledge to take off its sting, and leave us with a rather unexciting set of remarks about the hazards of sublunary reasoning? Well, that's exactly what I hope it is. Philosophical sanity doesn't have to be exciting. Sometimes quite the opposite is what is needed.[6]

What then do I know by reflection alone about, say, a putative case of universal instantiation, 'Everything must perish, therefore Turandot must perish'? What I *do* know is that, provided 'Turandot' refers, the inference is valid. I'm not supposed to know by reflection alone whether 'Turandot' refers: rather, when inferring I either know it on independent grounds or *riskily* take it for granted. Likewise with my staple example about the two dogs. I have talked elsewhere of *unsafe reasoning* (FARIA 2009), and the point I'm wanting to drive home here is that reasoning is bound to be always, to some extent, an unsafe business; also, that this is something we are apt to not understand about the subject matter of logic.

The basic idea here (minus the allusion to risk which matters to me) was encapsulated in Stephen Schiffer's response to Paul Boghossian's paper 'Externalism and Inference', back in 1992. Rejecting the suggestion that appeal to narrow content was mandatory in order to take into account the fact that the ostensibly faulty reasoning of slow-switched Peter made sense from his (Peter's) point of view, Schiffer

---

[6] 'We want to replace wild conjectures and explanations by quiet weighing of linguistic facts.' (WITTGENSTEIN 1967, § 447).

asked: 'Can't we explain Peter's being epistemically justified by observing that his belief was produced by reasoning of a form guaranteed to be valid but for undetectable externalist contingencies?' (SCHIFFER 1992: 37)

I think Schiffer's suggestion may be expanded into an account of our entitlement to take for granted, *as we all do most of the time*, that what he calls 'undetectable externalist contingencies' are not going on. We need such an account, at any rate, if we are to make good the claim that the external individuation of thought contents is no hindrance to our being able, more often than not, to know by reflection alone whether an inference is deductively valid.

Now it would be nice to have a principled way of sorting out the relevant kinds of cases here. At a bare minimum, it seems to me that two main sorts of cases should be distinguished: those in which a rational subject will be expected to *know* that the relevant empirical facts obtain, and those in which we are entitled to *take for granted* that the relevant empirical facts obtain.

The line is rather thin here, but I have in mind the contrast between what we are able to *come* to know (a proper subset of which will comprise facts we cannot afford *not* to know, the domain of epistemic obligation), and what lies beyond the reach of cognitive achievement. It is open to me to come to know that my neighbor is a breeder of Golden Retrievers, not that objects around me are usually stable and are not switched every time I blink. It is open to me to come to know that here is a hand and here is another, not that there is an external world.

If that distinction is on the right track, then a place must be secured, in the epistemology of reasoning, as in epistemology *tout court,* for the notion of inexcusable ignorance.[7] We will accordingly take into

---

[7] In *On Denoting,* Wittgenstein writes: 'That I am a man and not a woman can be verified, but if I were to say I was a woman, and then tried to explain the error by

account the possibility that, unlike what happens with the exiles on Twin Earth (as the story is usually told), a reasoner will often be in a position to acquire the relevant information about her circumstances and environment, and thereby raise and settle explicit questions about existence, permanence, identity or difference of the objects she reasons about; also, that failure to do that may be failure to comply with an epistemic duty. I may, as a matter of fact, disregard the possibility that the dog I saw on Tuesday is not the dog I saw on Friday, but that may be a simple case of epistemic negligence given antecedently available information.

But there is nothing I can do to ultimately satisfy myself that objects around me are usually stable and are not switched while I blink. So much has to be simply taken for granted, *even as it may, on occasion, turn out false.* Sometimes life imitates art: in 2009 two white mules artfully painted so as to look like zebras were on display at the Gaza Zoo Mahra Land. It is not known that the Zoo keepers ever read Fred Dretske. It is safe to say, on the other hand, that their contrivance failed to undo the philosophical verdict that, by and large, being a cleverly disguised mule is not a relevant alternative to being a zebra.

The same holds for our background assumptions about existence, permanence, identity, or causal properties of objects around us. To the extent that we are bound to take for granted (as opposed to come to know) that these assumptions hold, we may be said to be *a priori*, even though defeasibly, entitled to them.

Now whether we are so entitled would seem to depend, as my appeal to the notion of relevant alternatives was meant to suggest, on what is *normal* - on what is usually the case; on what Wittgenstein was wont to call *natural history*. And here an objection may be raised, to the effect that it is not up to philosophy to settle, from the armchair as it

---

saying I hadn't checked the statement, the explanation would not be accepted.' (WITTGENSTEIN 1969: § 79)

were, what is normal. Here's David Sosa again, commenting on Tyler Burge's claim that if there are any cases of equivocation due to slow switching (which Burge goes to great lengths to hold that there aren't), 'they are marginal' (BURGE 1998: 368). Writes Sosa: 'I find Burge's approach here troubling: one might have thought that whether such cases are typical or marginal - how often they occur, that kind of thing - is not for philosophers to judge.' (SOSA 2005: 224)

I want to resist that thought. I claim that we must be entitled to our background empirical assumptions *most of the time* if reasoning is to be possible *at all*.

Can I prove that? I think John Campbell showed how to do it.[8] Suppose you are holding a ball which feels soft and looks red. You think 'This ball is soft' and 'This ball is red'. Since you take both judgments to be true and don't even dream of their not being about the same object, you are ready to draw the conclusion 'Something is both soft and red'.

But what if, instead, you pause to do a bit of epistemology? Unleash your imagination; and suppose the visual image of a ball is being conveyed to you through a set of mirrors from a rigid red ball in the adjacent room while your tactile sensations are of a soft green ball which you are at present unable to see.

The supposition will be wild, as is often the case with what are called 'cases' in contemporary analytical philosophy; but notice that just *thinking* about it is enough to deprive you of your inferential innocence. For, having raised the possibility, you have *ipso facto* introduced two new ways of thinking about what you had previously thought of as 'This ball'. You now have 'The ball I'm touching' and 'The ball I'm seeing', and an identity problem.

In Campbell's lingo, you ceased to *trade upon* the identity of the seen and the touched object; instead, you introduced a divide in the

---

[8] See, again, CAMPBELL 1987.

information coming to you from the object, so that your premises are now of the form 'a is F', 'b is G'. And now, of course, it will be wise of you to refrain from drawing any conclusion from this couple of logically independent judgments.

What I want to stress is that *this was your own deed*: no Evil Genius needs to have pulled the old switcheroo so as to have you take this for that; no overnight travel from Earth to Twin Earth must have taken place. You have, by your own unassisted means, undone whatever entitlement you may have had to draw the conclusion.

Compare this to the symmetrically opposite case: there actually are two objects, a rigid red ball and a soft green ball, but you cannot track the fact that your information comes from these two distinct sources, so that you keep thinking of both as 'This ball'. *Now* you are prone to draw a false conclusion from a set of true premises, exactly like those unlucky exiles on Twin Eart, Boghossian's Peter and his kin.

The conclusion I'm aiming at is now at hand. In section 401 of *On Certainty* (1969), Wittgenstein writes: 'I want to say: propositions of the form of empirical propositions, and not only propositions of logic, form the foundation of all operating with thoughts (with language).'

I think that is a difficult idea, not because it is recondite, but because it goes against the grain of a deeply entrenched picture of what Wittgenstein here describes as 'operating with thoughts'. On that picture, thought evolves in a frictionless medium, sealed off from the vagaries of actions, the hazards of causation and the accidents of history.

To see what exactly Wittgenstein is opposing to that picture, though, we need to clarify what he means by the *form* of an empirical proposition, as it is by no means obvious (to say the least) what the word 'form' is doing there. This is as vital to reach a satisfying reading of section 401 as it is to understand why Wittgenstein writes, in the remark which immediately follows it: 'In this remark the expression

"propositions with the form of empirical propositions" is itself thoroughly bad' (WITTGENSTEIN 1969: § 402).

The phrase surfaces rather early in *On Certainty,* in the context of introduction of the celebrated image of 'the river-bed of thoughts'. Here is section 96: 'It might be imagined that some propositions, of the form of empirical propositions, were hardened and functioned as channel for such empirical propositions as were not hardened but fluid; and that this relation altered with time, in that fluid propositions hardened, and hard ones became fluid.'

Then in section 308, with respect to what he will call (from section 341 onwards) 'hinge' propositions, Wittgenstein writes: '[...] we are interested in the fact that about certain empirical propositions no doubt can exist if making judgments is to be possible at all. Or again: I am inclined to believe that not everything that has the form of an empirical proposition *is* one.' Here 'having the form of an empirical proposition' is contrasted with *being* an empirical proposition, and it is, as it were, as an afterthought ('or again') that the amendment is recorded.

Now what is the form of an empirical proposition? Come to think, what is the *form* of a proposition?

Wittgenstein's only explicit explanation of that notion[9] is apt to strike us as being as unoriginal as it is unpromising. It surfaces in the opening of 'Some Remarks on Logical Form': 'Every proposition has a content and a form. We get the picture of the pure form if we abstract from the meaning of the single words, or symbols (so far as they have independent meanings). That is to say, if we substitute variables for the constants of the proposition.' (WITTGENSTEIN 1929: 29). (Russell might have written that. He actually wrote it more than once, albeit in a variety of different wordings).

---

[9] To the best of my knowledge, it goes without saying.

But then the main thrust of 'Some Remarks on Logical Form' was precisely to expose the insufficiency of that conception of form when it is supposed to account for the fact that certain propositions 'stand fast to us' - even as they don't display the form (as defined) of a tautology. We know *a priori*, to take the most notorious example, that nothing can be red and green all over.[10] Likewise: 'One shade of colour cannot simultaneously have two degrees of brightness, a tone not two different strengths, etc. And the important point here is that these remarks do not express an experience but are in some sense tautologies.' (WITTGENSTEIN 1929: 32)

In some sense: not, to be sure, in the sense of the *Tractatus* - cf. in particular the program for an analysis of (what Russell had called) synthetic incompatibilities announced in WITTGENSTEIN 1921, section 6.3751: 'It is clear that the logical product of two elementary propositions can neither be a tautology nor a contradiction. The statement that a point in the visual field has two different colours at the same time is a contradiction.' We then find Wittgenstein writing in 1929 that 'The mutual exclusion of unanalyzable statements of degree contradicts an opinion which was published by me several years ago, and which necessitated that atomic propositions could not exclude one another. I deliberately say "exclude" and not "contradict", for there is a difference between these two notions, and atomic proposition, although they cannot contradict, may exclude one another.' (WITTGENSTEIN 1929: 33)

I take the latter remark to be, historically, the first gesture at the idea of what Wittgenstein was later to call a *grammatical* proposition. *Grammatical*, not *logical* - but this, I submit, is just an interim choice of vocabulary. In a remark written in 1940, Wittgenstein explains: 'Sometimes an expression has to be withdrawn from language and sent for cleaning, - then it can be put back in circulation.' (WITTGENSTEIN

---

[10] Compare Russell's account of 'synthetic incompatibilities' in RUSSELL 1900: 17-24 and then in RUSSELL 1903, chapter XXVII ('Difference of Sense and Difference of Sign'): 227-233

1977: 39). Now I take it that the most important family of expressions withdrawn for circulation from 1929 onwards is 'logic', 'logical', 'logically'.[11] Yet they reappear *en masse*, replacing 'grammar' and 'grammatical', in the last writings, most notably in *On Certainty*.[12] At that point, I surmise, the cleaning was completed, and the words ready for reentering in circulation.

In the meantime we were introduced, time and again, to cases of propositions whose form (as defined in 'Some Remarks') disguises their grammatical (that is, their *logical*) role. Hence we find. e.g., in what was formerly known as Part II of the *Philosophical Investigations* the remark that 'War is war' is not an example of the law of identity (Wittgenstein 1953b: § 311). 'War is war' is, in other words, an empirical proposition masquerading as a logical one - an empirical proposition with the form of a logical one.

Now one thing I am *not* claiming is that the view that the use of a proposition (its role in a language game) is what determines its logical form - that this view is a novelty of Wittgenstein's later philosophy. Here is *Tractatus* (WITTGENSTEIN 1921) 3.327: 'A sign does not determine a logical form unless it is taken together with its logico-syntactical employment.[13]

We are now in a position to read section 402 of *On Certainty* in its entirety:

---

[11] They are used, to be sure, in the *Philosophical Investigations*, but only in the context of Wittgenstein's reappraisal of the philosophy of the *Tractatus*, most notably in sections 89-114, 242 and 437.

[12] See, e.g., the remarks numbered 21, 26, 36, 43, 48, 52, 53, 56, 59, 68, 82, 98, 110, 155, 194, 308, 319, 342, 353, 375, 401, 447, 454, 475, 501, 618 and 628 of WITTGENSTEIN 1969. By contrast, 'grammar' features on 51, 313 and 433 - and, as far as I could check, that is it.

[13] Of course there is leeway here for disagreement about the Tractarian notion of the logico-syntactical employment (*Verwendung*) of a sign. I do not have the space to discuss the issue in this paper. Let me just note that, here as elsewhere, the continuity between the *Tractatus* and Wittgenstein's later work is way greater that it was (is?) usually taken to be. See, as especially relevant to the issue, *On Certainty* § 321.

> In this remark the expression "propositions with the form of empirical propositions" is itself thoroughly bad; the statements in question are statements about material objects. And they do not serve as foundations in the same way as hypotheses which, if they turn out to be false, are replaced by others.
>
> ... und schreib getröst
>   "Im Anfang war die Tat."

That the statements in question are statements about material objects is here meant to *explain* what was wrong with the phrase 'the form of an empirical proposition'. The statements in question are not, in other words, grammatical (that is, *logical*) propositions masquerading as empirical ones (the dual of the case discussed in Wittgenstein 1953b § 311, on 'War is war'). They are *themselves* empirical propositions - which, however, do not play the role of hypotheses open to refutation by experience: precisely like the empirical assumptions (about existence, permanence, identity, or causal properties of objects) which I claimed are *not* hidden premises in enthymematic reasoning, playing rather the role of background assumptions in any assessment of the logical validity of inferences:

> Here one must realize that complete absence of doubt at some point, even where we would say that 'legitimate' doubt can exist, need not falsify a language game. For there is also *something* like another arithmetic.
>
> I believe that this admission must underlie any understanding of logic. (WITTGENSTEIN 1969 § 375)

We now turn full circle: if there was something that David Kaplan 'was not understanding' about 'the subject matter of logic', the something was precisely the role played by empirical assumptions in logical appraisal - a role whose recognition 'must underlie any understanding of logic'.

And, again, it is only in its logico-syntactical employment that a sign determines a logical form: what *we do* with a propositional sign, in other words, is what makes it either a logical or an empirical proposition - and, if the latter, either a hypothesis or a 'hinge' on which hypotheses turn. So 'write with confidence: "In the beginning was the deed".'

I talked of a picture of thought as evolving in a frictionless medium, sealed off from 'the vagaries of action' - a picture whose complement is the picture of logic as concerned exclusively with 'the verities of meaning'. That pictures holds philosophers after Kaplan's heart (and their name is Legion) captive,[14] even as they are cognizant, as Kaplan himself is, that something may be very wrong about it. For, let there be no doubt, it is a *very* powerful, mesmerizing picture.

Once, however, we have reconciled ourselves with the couple of platitudes I've been insisting upon, things begin to look otherwise: and then, just possibly, the idea that the external individuation of thought contents is bound to make the validity of our inferences imperfectly accessible will cease to seem to call for either a *reductio* of anti-individualism or for some version of "compatibilism". We may just try and get used to the idea that, as Timothy Williamson puts it with respect to this very topic, 'real life is messy' (WILLIAMSON 2000: 16).

Here is Robert Louis Stevenson, and let him have the last word: 'I told him I was not much afraid of such accidents; and at any rate judged it unwise to dwell upon alarms or consider small perils in the arrangement of life. Life itself, I submitted, was a far too risky business as a whole to make each additional particular of danger worth regard.' (STEVENSON 1879: 121)

Universidade Federal do Rio Grande do Sul, Brazil

paulo.faria@ufrgs.br

REFERENCES

BOGHOSSIAN, Paul A. (1992). 'Externalism and Inference', *Philosophical Issues* **2**: 11-28

BOGHOSSIAN, Paul A. (1884). 'The Transparency of Mental Content', *Philosophical Perspectives* 8: 33-50, reprinted in *Content and Justification* (Oxford: Clarendon Press, 2008): 159-176

BURGE, Tyler (1988) 'Individualism and Self-Knowledge', *Journal of Philosophy* **85**: 649-663, reprinted in LUDLOW & MARTIN (1998): 111-127

---

[14] Cf. Wittgenstein 1953a § 115.

Burge, Tyler (1998). 'Memory and Self-Knowledge', in Ludlow & Martin (1998): 351-370

Camp, Jr., Joseph L. *Confusion: a Study in the Theory of Knowledge* (Cambridge, MA: Harvard University Press, 2002)

Campbell, John (1987). 'Is Sense Transparent?, *Proceedings of the Aristotelian Society* **88** (1987/1988): 273-92

Campbell, John (1994). *Past, Space, and Self* (Cambridge, MA: The MIT Press).

Collins, John, (2008). 'Content Externalism and Brute Logical Error', *Canadian Journal of Philosophy* **38**: 549-574

Faria, Paulo (2009). 'Unsafe Reasoning: a Survey', *Dois Pontos* **6**: 185-201

Kaplan, David (1989) 'Afterthoughts', in J. Almog, J. Perry & H. Wettstein (eds.), Themes from Kaplan (Oxford: Oxford University Press)

Ludlow, Peter (1998) 'Social Externalism and Memory: a Problem?', *Acta Analytica* **14**, reprinted in Ludlow & Martin (1998): 311-317

Ludlow, Peter (2004). 'What Was I Thinking? Social Externalism and Shifting Memory Targets', in R. Shantz (ed.), *The Externalist Challenge* (Berlin: De Gruyter): 419-426

Ludlow, Peter & Martin, Norah, eds. (1998). *Externalism and Self-Knowledge* (Stanford: CSLI Publications)

Recanati, François (2012). *Mental Files* (Oxford: Oxford University Press)

Russell, Bertrand (1900). *A Critical Exposition of the Philosophy of Leibniz* (London: Routledge, 1992)

Russell, Bertrand (1903). *The Principles of Mathematics* (London: George Allen & Unwin, 1964)

Schiffer, Stephen (1992). 'Boghossian on Externalism and Inference', *Philosophical Issues* **2**: 29-38

Sorensen, Roy A. (1998). 'Logical Luck', *The Philosophical Quarterly* **48**: 319-334

Sosa, David (2005). 'The Inference that Leaves Something to Chance', in Sanford Goldberg (ed.), *Internalism and Externalism in Semantics and Epistemology* (Oxford: Oxford University Press): 219-234

Stevenson, Robert Louis (1879). 'Travels with a Donkey in the Cevennes', in *An Island Voyage, including "Travels with a Donkey"* (New York: Cosimo: 2007)

Williams, Bernard (1976). 'Moral Luck', *Proceedings of the Aristotelian Society*, supplementary volume 50, reprinted in *Moral Luck* (Cambridge: Cambridge University Press): 20-39

Williamson, Timothy (2000). *Knowledge and its Limits* (Oxford: Oxford University Press)

Wittgenstein, Ludwig (1921). *Tractatus Logico-Philosophicus*, tr. David Pears & Brian McGuinness (London: Routledge, 1961)

Wittgenstein, Ludwig (1929). 'Some Remarks on Logical Form', *Proceedings of the Aristotelian Society, Supplementary Voltume* **9**: 162-171. reprinted in *Philosophical Occasions 1921-1951*, ed. James Klagge & Alfred Nordmann (Indianapolis & Cambridge: Hackett Publishing Company, 1993): 29-35

WITTGENSTEIN, Ludwig (1953a). *Philosophical Investigations*, revised 4th edition by P. M. S. Hacker & Joachim Schulte (Oxford: Blackwell, 2009)

WITTGENSTEIN, Ludwig (1953b). 'Philosophy of Psychology: a Fragment' in WITTGENSTEIN 1953a: 185-243

WITTGENSTEIN, Ludwig (1967). *Zettel*, ed. G. E. M. Anscombe & G. H. von Wright, translated by G. E. M. Anscombe, 40th anniversary edition (Berkeley & Los Angeles, CA: University of California Press, 2007).

WITTGENSTEIN, Ludwig (1969). *On Certainty*, ed. G. E. M. Anscombe & G. H. von Wright, translated by Denis Paul & G. E. M. Anscombe (Oxford: Blackwell)

WITTGENSTEIN, Ludwig (1977) *Culture and Value / Vermischte Bemerkungen*, ed. G. H. von Wright & Heikki Nyman, tr. Peter Winch (Chicago: The Universitty of Chicago Press, 1980)