What if *all* truth is context-dependent?

Bruce Edmonds

Centre for Policy Modelling,
Manchester Metropolitan University,
Aytoun Building, Aytoun Street, Manchester, M1 3GH, UK.
http://www.cpm.mmu.ac.uk/~bruce

Abstract. This paper argues that truth is by nature context-dependent – that no truth can be applied regardless of context. I call this "strong contextualism". Some objections to this are considered and rejected, principally: that there are universal truths given to us by physics, logic and mathematics; and that claiming "no truths are universal" is self-defeating. Two "models" of truth are suggested to indicate that strong contextualism is coherent. It is suggested that some of the utility of the "universal framework" can be recovered via a more limited "third person viewpoint". **Keywords:** philosophy, universality, context, truth, knowledge.

1. Introduction – the position of strong contextualism

The standard philosophical world-view presumes that there *exist* universal truths – propositions that hold regardless of context. In other words, that truth *itself* is somehow separate (or separable) from the messy contingencies of our world. Furthermore it often assumes that sometimes one can *know* these truths. This paper considers the opposite situation, namely, the possibility that truth is by its nature context-dependent – that truths only have *meaning* in a limited set of contexts and thus they are only applicable in those contexts. One consequence of this position is that for every proposition there is a context in which it does not hold.

This can be seen as the result of an ineradicably association of a truth with its development, expression and application. That although *some* abstraction and generalisation away from particularities is possible, it is not possible to abstract away from all contexts and generalise to complete universality. For example, "red is a colour" may be a generalisation from several many context-specific experiences of different shades of red, but if one abstracts away from all experience contexts it loses its meaning and there will be some contexts where it is simply inapplicable. This includes those in categories such as "synthetic", "logical" and "a priori". As I explain in section 5, I use the work "truth" because that is the most sensible and intuitive term, and I am arguing against the distinctions implicit in the expression/proposition and truth/knowledge divides – I do mean that "truth" itself is context-dependent and not merely knowledge or expression.

This goes beyond saying (with Penco, 1999) that "there is no ultimate outer context" for as Roger Young (1999) has shown it is possible to define absolute truth

in such a situation using an abstract notion of quantifying across ever-wider contexts – something can be defined as absolutely true if it is true in all contexts containing the one it is posited in. Such a quantification is, of course, an absolute notion – standing above a never-ending progression of ever wider contexts. Quite apart from this being a utterly impractical process (how does one ever know that there will not be another more general context in which a particular truth is false?), this notion has already assumed the possibility of such universal notions as quantification over contexts in order to establish the possibility of defining absolute truth in a world where there is no "outer context". It is possible that this sort of notion can be applied reflectively to itself and the further quantification over meta-contexts that would be necessary etc. but then what grounds would one have for saying that the result contains absolute truths? Each level of establishing truth involving a quantification over contexts is limited to its infinite set of contexts – it is conceivable that such a trans-infinite recursion results in truly universal truths, but there are no firm grounds for saying this would be the case (even if it were possible).

I am going further than saying there might be an infinite regression of wider-contexts, I am looking at the possibility that any particular truth is linked to specific contexts as part of its nature and cannot be applied regardless of context. This paper argues for the coherency of such a position, looks at some of the reasons why one might suppose this is the case and suggests how one might recover almost all of the utility of the fictional "absolute frame of reference" as a means of facilitating discussion by introducing "the philosophical context" without some of the more misleading aspects implicit in the absolute version.

For want of a better term I will call such my position "strong contextualism". This is distinguished from both relativism and, what I will call "weak contextualism" – that there may be universal truths but we never *know* them (in the strong philosophical sense of believing them with good justification). Relativism is that the truth of all propositions is relative to the person or viewpoint. It allows for the possibility that from within a view point a proposition might be held (with good reason) to be true over all contexts, unlike strong contextualism. Also strong contextualism allows for the possibility of objective knowledge (unlike many forms of relativism) because a proposition can be simultaneously context-dependent and objective – as long as the context in which a proposition holds can be reliably identified by third persons (regardless of their viewpoint) its truth within that context can be independently tested.

This paper has a lot in common with the conclusion of (Penco, 1999), but whereas he approaches his goal subtly examining some of the problems with objective approaches to context, I am aiming to tackle the question "head on".

The rest of the paper will be divided as follows. Section 2 will briefly recap my previous analysis of the nature of context. Sections 3, 4, 5 and 6 will deal with some obvious objections to the position: that there appear to be some universal truths; that strong contextualism is self abnegating; that I am simply confusing truth with knowledge or belief and that context-dependent truths can be simply converted to context-independent ones. Sections 7, 8 and 9 will look at some more positive arguments: suggesting some models to show the coherency of the position; looking at some positive reasons why it is sensible; and showing that the utility of the "absolute framework" as a means of facilitating discussion can be recovered within strong

contextualism. Section 10 will indicate the context of the paper and I will conclude in section 11.

2.   Context

The following analysis of context follows (Edmonds, 1999) which goes into it in much greater depth.

   The possibility of learning and inference in our complex world is dependent on the fact that many of the possible causes of events remain relatively constant in most circumstances. If this were not the case we would need to include all the possible causes in our models. This relative constancy is what makes knowledge possible: we can learn a model in one circumstance and apply it in another circumstance that is sufficiently similar to the first. The label of 'context' is as a stand-in for those factors that are not explicitly included in the models we learn, or, to put it positively, those factors that we use to recognise when a model is applicable.

   It is the possibility of the transference of knowledge from the circumstances where they are learnt to the circumstances where they are applied which allows the emergence of context. The utility of 'context' comes from the possibility of such transference. If this were not feasible then 'context', as such, would not arise. Context and knowledge are contingently possible because of the nature of the world. This is illustrated below in figure1.
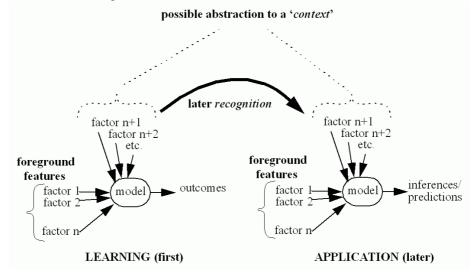

Fig. 1. Context in the transference of knowledge between learning and application

   For such a transference to be possible it is necessary that:
   • some of the possible factors influencing an outcome are separable in a practical way;
   • a useful distinction can be made between those factors that can be categorised as foreground features (including 'causes') and the others;

- the background factors are capable of being recognised later on;
- the world is regular enough for such models to be at all learnable;
- the world is regular enough for such learnt models to be at all useful when applied in situations where the context can be recognised.

While this transference of knowledge to applicable situations is the basic process, observers and analysts of this process might identify some of these combinations of features that allow recognition and abstract them as a 'context'. This usually is possible because the transference of knowledge as models requires that the agent doing the transference can recognise these characteristic combinations, so it is possible that an observer might also be able to do so and give these combinations names. On the other hand the underlying recognition mechanism may be obscure. Of course, it may be that the agent doing the transference itself analyses and abstracts these features, and thus makes this abstract available for reflective thought.

Specifying a context by listing *all* the conditions/causal factors that are not explicit in our knowledge is impossible – this is akin to listing all the things that are *not* in a box, the list is infinite. This is why contexts are things that are primarily *recognised* rather than inferred. A context independent truth, from this bottom-up perspective, would have to combine the explicit content of a model with the implicit content of the context. This is the reason why truth is so dependent upon context – because a large portion of its *content* necessarily resides there.

In this picture of things (a picture which is rooted in the process by which truths are actually established): beliefs are learnt by interaction with the world and others; knowledge and belief are distinguished by what can be independently validated as reliable; and truth is the content of this knowledge. This is opposite to a picture whereby (at least some) truths can be established by pure argument and only applied with the addition of contingent detail later. Such an idealist picture has big problems defining knowledge, because from this (theoretically) absolute standpoint there is no necessity for any connection between our belief, our reasons for having the belief and the truths it has to coincide with it is to obtain to knowledge.

3.   Apparently universal truths

The first objection I will deal with is that fact that there appear to be universal truths, for example those given to us by logic, mathematics and the "hard" sciences. While I could not hope to consider the universality of all such truths in one paper, I do hope to indicate via a few archetypal examples how these truths might have "hidden" contexts in which their truth is firmly embedded.

3.1   Physics

The archetypal "universal truths" of physics are Newton's laws of motion. These are constant across frames of reference that are travelling at a constant velocity relative to each other. It is, of course, now well known that there *are* some conditions

of application for these laws – they hold only in the macroscopic world for velocities that are small relative to that of light. Now it may be supposed that the process of the contextualisation of Newton's laws have reached an end with quantum corrections and extensions due to the Theory of Special Relativity, but this is not the case. General relativity means that another condition is that you are not near a large mass (which produces an effect indistinguishable from a large acceleration) and more recently the speed necessary for relativistic effects to become apparent has come sharply down in special circumstances where the speed of light has been slowed (to lower than 1 m/s).

In fact there has been a continual stream of thinking since Einstein, postulating even more radical contractions of universality. There is a school of thought that the different types of forces (and their accompanying "laws") only separated as the universe cooled in the moments after the Big Bang (Weinberg, 1988). This would mean that many of the laws of physics depend upon a relatively cool context for their existence. In another direction the "Anthropic principle" (Barrow and Tippler, 1986) suggests that many aspects of our universe can be "explained" because if they were not so we would not have evolved to formulate the laws. Quantum physics suggests that new universes might be being created all the time forming a "foam" of distinct universes which might have different laws of physics – it is speculated that only those sufficiently stable to support intelligent life will have its laws recognised and formulated.

Philosophers such as Nancy Cartwright (19983) and Richard Giere (1988) who study the process of science have documented how the application of laws to the world is not a neat, axiomatic one but grounded in a rich scientific context which provides the all-important rules as to how one relates models to situations. Training is required in order to be able to successfully apply the "laws of physics" because there is no infallible rule book but rather a mixture of processes is required including fuzzy recognition. Nancy Cartwright goes as far as distinguishing theoretical and phenomenological laws, claiming that theoretical laws are strictly false whilst phenomenal laws are specific to a particular situation.

Thus although a naïve[1] picture of physics characterises it as universal, the reality (especially as revealed in its practice) shows otherwise. This is not overly surprising – the training of a physicist is long and is not merely concerned with how to look-up various laws. Rather the efficacy of physics lies in the ability to recognise a rich set of contexts[2], to choose and then adapt the relevant techniques.

3.2   Logic

Another archetypal truth is the inference pattern implicit in "Aristotle is a man, all men are mortal, therefore Aristotle is mortal". This seems to be universal, but this is merely a limitation upon our imagination – consider "This is a sentence, the meaning

---

[1] Either naïve or radically systematised and simplified for didactic purposes in, for example, textbooks.

[2] As Cartwright documents scientists frequently apply different sorts of model with different sorts of approximation simultaneously, even when these are (in strict terms) incompatible.

of sentences are determined by their use, so this sentence's meaning is determined by its use" (i.e. *this* use)[3]. The reasoning now no longer follows so inexorably – a pattern designed for exterior use may fall down in the presence of self-reference.

The standard response by logicians is that there is nothing wrong with the inference pattern itself, it has been merely applied wrongly but, of course, they are unable to give criteria for the correct application of inference patterns except by reference to specific applications where it gives the answer they expect. They push the problem into someone else's court in the hope of retaining universal validity[4]. This attempt is undermined by the proliferation of alternative logical formalisms, each different from the others. They can not all be universally correct in themselves, and although there are still a few logicians who are still holding out for the "one true logic", the rest are forced by the practicalities into accepting that one chooses one's logic according to the context.

Practical reasoning – that is reasoning that actually gives useful conclusions that *work* – is a world away from its models in formal logic and inevitably makes heavy recognition and use of context. This is illustrated by the following classic example; when subjects were told "If the light is red then the car stops" and asked "What can you deduce from the fact that the car stops?", the majority (of people who are not formally trained in logic) reply "The light is red". There was some discussion on the PHIL-LOGIC (philosophical logic) mailing list as to why the subjects made the "wrong inference", with the most popular explanation being that implication (if) was confused with equivalence (if and only if). Whereas a more believable explanation is that the question implicitly indicated to the subjects that the appropriate context is "stopping at traffic lights" where the inference is correct (there not being any other cause of stopping).

Again what looks like universality from afar is necessarily context-dependent in use – so called universal laws of logic are not universally applicable.

3.3 Mathematics

Pure mathematics aspires to a world of its own. It is concerned with what can be formally proven given certain structures, assumptions etc. For example, given Peano's axioms for arithmetic, the standard notation and some standard logical inference operations, one can prove the statement "1+1=2". Does this not mean that that "1+1=2" is a universal truth, devoid of context? I would argue not.

There are two interpretations of "1+1=2": that it is a formal sequence of symbols that are provable from other sequences using formal rules; or that it expresses a fact about objects in the world, namely, that putting one object together with another object gives you two objects. I will consider these one at a time.

---

[3] Or more accurately: "This is a sentence, Wittgenstein's later philosophy suggests that the meaning of sentences are determined by their use, so Wittgenstein's later philosophy suggests that the meaning of this sentence is determined by its use".

[4] A notable exception to this is (Adams and Levine 1976), who start to investigate the conditions under which logical inference will hold in the face of uncertainty.

If one takes the formal interpretation, then the statement "1+1=2" doesn't have any *meaning* outside that given by the formal system it is part of – it derives *all* of its truth, meaning and relation to other statements from that system. It is, in other words, entirely dependent upon the context of that system.

If "1+1=2" is about the world, it is not making an empty statement – it is saying, for example, that two objects retain their identity when considered together, that they don't merge and become one ("1+1=1"), or even disappear ("1+1=0"). Now we know that, in most circumstances this is a sensible way to consider the world – it impinges upon us as separate and identifiable chunks. There is evidence from child psychologists that we have, at a very early age, an ability (or bias) to view the world in such units. The "countability" of objects is such a pervasive part of our experience that we forget upon the huge assumptions and properties upon which it is based. Such an interpretation does not survive a move to the context of the sub-atomic world; there it may make far more sense to not to consider particles as discrete but as observable manifestations of a single, continuous wave function. The point is that when treated as a statement about the world, it is as context-dependent as any other such statement. It does not lose this context-dependency just because it happens to be expressible in a formal system.

There is a third interpretation, of course, but this is one that already *presumes* universality as its starting point. It is that "1+1=2" is somehow indicative of a universal truth about the relation of *oneness* and *twoness* that can be seen as an abstraction of all the real world interpretations of "1+1=2". The existence of such a truth is entirely by presumption, claims for its universality are not based on any evidence or argument but seem to rest mainly on the fact that it can be conceived to be so[5].

The importance of mathematics comes from the fact that one can establish strong mappings from it to aspects of the world and its utility from moving between the formal realm where syntactic moves are made (for example in a calculator) and what we are considering (for example sheep). It is when a mapping is established that the power of mathematics becomes manifest. I suspect that it is this immensely useful ability to map between contexts that has led some to make the mistake that mathematical truths are universal.

4.   The coherency of strong contextualism

The next obvious objection to strong contextualism is that it is self-defeating. The argument goes something like this: if all truths are context-dependent so is this, therefore there is a context in which this is not true and so, in general, strong contextualism is false. This argument is similar in structure to many paradoxes including the liar paradox and Russell's set-theory paradox.

Of course, the nub of the argument is in the words "in general". I *do* claim that truth has no meaning *in general*, and that includes the truth of strong contextualism. However, that does not mean that there *are* universal truths – to get to the *existence* of

---

[5] Such an idea can legitimately be used as a useful "fiction" or "short-hand" to simplify expression when the context is well understood, but this is very different from believing it!

universal truths from this pseudo-paradox one would have to do something such as going from $\neg\forall t[\exists c_1(\text{ist}(t,c_1)) \rightarrow \exists c_2(\neg\text{ist}(t,c_2))]$ – a denial of the statement that every statement true in some context is false in another – to $\exists t\forall c(\text{ist}(t,c))$ – the existence of a statement true in any context – (here t is a truth, c_1 and c_2 are contexts, $\forall t$ stands for "all truths", $\exists c_1$ stands for "there is a context", etc.) as in the following reasoning:

```
¬∀t[∃c₁(ist(t,c₁)) → ∃c₂(¬ist(t,c₂))]    assumption
∃t¬[∃c₁(ist(t,c₁)) → ∃c₂(¬ist(t,c₂))]    ¬∀ to ∃¬
∃t¬[¬∃c₁(ist(t,c₁)) ∨ ∃c₂(¬ist(t,c₂))]    defn implication
∃t[∃c₁(ist(t,c₁)) ∧ ¬∃c₂(¬ist(t,c₂))]    de Morgan
∃t[¬∃c₂(¬ist(t,c₂))]                      conjunction out
∃t[∀c₂¬(¬ist(t,c₂))]                      ¬∃ to ∀¬
∃t[∀c₂(ist(t,c₂))]                        double negation
∃t∀c(ist(t,c))                           renaming symbols
```

The trouble is that to push this argument through one *already assumes* the existence of universal truths, including: the ability to quantify over all truths and all contexts, and that negation and implication are classical (which implicitly utilises an absolute framework in its assumptions, for example the law of the excluded middle). Basically one is using a highly non-constructive proof that presumes that things such as the law of the excluded middle works when quantifying over *all* truths and sets. Such quantification involves quantification over structures bigger than the classes that caused naïve set theory such problems (e.g. the set of all sets), since "ϕ is a set" is a truth for all sets ϕ.

Strong contextualism *is* context-dependent and thus is not as strong as a putative universal truth would be (if one existed in any meaningful sense), but it could well be as strong as other truths that we rely upon, such as: Schrödinger's Wave equation, the logical rule of Modus Ponens or the arithmetical statement "1+1=2". The fact that strong contextualism is false *in general* is unthreatening, because it is trivial.

What happens if we try the argument the other way around: if in any meaningful sense *all* truths are context-dependent, then this (or some closely related statement such as "all truths *except this one* are context-dependent") is universally true. Again this fails because it assumes that just because a statement has the word "all" in it that this must be a (truly) universal quantification rather than some *meaningful* quantification relevant to the context. This brings up the nature of the context of strong contextualism itself which I will discuss in section 7 below.

Strong contextualism is only self-negating if one takes tries to examine it from a universal stance – the moral of these arguments is that one cannot safely mix the two. In the next section I look at some more positive ways of addressing the coherency of strong contextualism.

5. "Knowledge" vs. "Truth"

The third objection is simply that I am using the wrong names: what I mean by "truth" is what others call "knowledge" (or even "belief"). That "truth" can't be context-dependent by definition, because then it wouldn't be "true"[6]. To distinguish these I shall call the "true"="true independent of context" nomenclature "universalist nomenclature" and the usage as in most of this paper "contextualist nomenclature". Using universalist nomenclature the position of this paper would be that *there are no truths*. Put this way the thesis sounds idiotic but this is due to the fact that philosophy routinely sets an *unrealistically* high standard for truth (namely that it should be true *without exception* in *any* context) whilst at the same time purporting to be able to say something about them.

There are problems with the universalist nomenclature. *Firstly*, it produces strongly counter-intuitive (dare I say *unreal*) results. According to this it is not "true" that "the 192 bus goes to Manchester" despite the fact that it does! This is due to the fact that it is possible to conceive of contexts where it is not true (e.g. Manchester in the U.S.). In other words universalist nomenclature diverges sharply from common parlance in, for example, most of the sciences. *Secondly*, it *presupposes* that truth and knowledge are separable, with truth being the universal ideal which we can sometimes obtain. This is precisely what I am arguing against.

So to be clear, I arguing that statements that almost everybody would, with good reason, take as true are context-dependent. This includes such as "1+1=2", "bachelors are unmarried", "Julius Caesar conquered Gaul" and the laws of physics. I choose to call these "truths" because they are statements that are true! I am claiming that even these sorts of things are context-dependent. If the reader wishes to go through the paper replacing all instances of "true" and "truths" with "known" and "knowledge" they are welcome to do so providing they then keep in mind that I am arguing (in this ridiculously strong universalist nomenclature) that *there are no truths* - or even *all statements are false* (again in this extremely strict sense).

6. Converting context-dependent to context-independent

The final objection I will deal with is this: that every context-dependent statement can be converted to a context-independent one by expressing its conditions explicitly. This could be expressed as moving from *ist(c, a)* to $c \rightarrow a$.

The reason why this is not completely possible is that whenever one context is a generalisation of some other contexts (e.g. the context of a "formal occasion", generalised from funerals, interviews, award ceremonies etc.) one inevitably looses some of the meaning in the less abstract contexts, because *meaning* ultimately derives from direct experience *in situ*. Thus if one continues to abstract away to more and more abstract contexts, one is left with literally *meaningless* expressions.

Contexts are always, to some extent, recognised rather than inferred. There is no set of conditions that one lay down that precisely describes the context. Seen as this,

---

[6] Though how one could go about *defining* "truth" without begging the question is beyond me!

basic context describe the edge of reason, or put more positively, *allow* the existence of "crisp" inference by prescribing a huge raft of "the given".

As anyone who has tried it will testify, making all the conditions explicit is utterly impractical, if not impossible.

7.   Two models indicating the possibility of strong contextualism

In this section I discuss some formal systems that might illustrate what truth under strong contextualism *might* be like. These analogies establish a sort of coherency for the position.

7.1   A Gödelian cascade of proof systems

Unlike truth, formal proof has always being accepted as system-dependent. A proof has to be incrementally constructed from the axioms and rules of the formal system – there is no such thing as a universal proof. The idea is to use an analogy where truth corresponds to a formally provable fact and contexts correspond to a formal system.

Consider a formal axiomatic system, sufficiently expressive to encode arithmetic, call it S_1. Gödel's incompleteness theorems tell us that there are facts in S_1 that are not provable in S_1. However it is possible to construct a meta-system S_2 in which these facts are provable. Gödel's incompleteness theorems apply equally to S_2 so there are facts about S_2 that are only provable in a meta-meta system S_3. In this way we can construct new layers S_n as required. All facts are provable in some system S_n but there isn't any system in which all facts are provable.

Truth may be like provable facts in these systems – all truths are dependent upon some context but this does not mean that there is a context in which all truths hold.

7.2   Alternative inner models in set theory

In the last analogy truths collect as you get further out, a truth dependent upon a lower context means that it also holds in all outer-contexts. Truth is not necessarily like that, so I give a second analogy where this is not true. This time the analogy is taken from set theory, where truth corresponds with the consistency of sets of axioms and context corresponds with inner models that show their consistency.

Now for any proper subset of the axioms of set theory (except foundation, which is there for technical reasons) there is a set in which all of this subset is true. This set is called an inner model because it is a model of the working of the subset of axioms. However (due to Gödel again) it is known that there can't be an inner model that showed the consistency of *all* the axioms of set theory, since that would mean that set theory proved its own consistency. Thus we have different inner models showing the consistency of different proper subsets of the axioms, but no one showing their consistency all together.

If truth does act as this analogy suggests, then there may well be different contexts for different truths but not one context for all truths.

8.   Positive reasons for supposing strong contextualism

The most obvious reason for supposing strong contextualism is that all truth *is* developed, established and applied in a context. It is pure supposition to believe that despite its "home grown" nature in practice that somehow there are truths that happen to be completely universal. Cases such as Aristotle's syllogism and "laws" of physics above may appear to suggest that they are universal but on closer inspection this is not necessarily the case.

Each truth having (at least one) context upon which it was dependent provides us with useful properties: a source for meaning; information to aid troubleshooting; and meta-knowledge about truth. I briefly look at these in turn.

A context can help provide the *meaning* of the truth by reference to the contexts of the truth's development, establishment and application. At the most basic level one learns the meaning of pain in a context where one experiences it, where one watches others experiencing it and when one tries to avoid it. The experience of other, somewhat incidental, facts that hold in these contexts gives the truth its "flavour", without which it would be difficult to relate to.

Knowing the context that a fact depends upon gives one valuable information about what assumptions to question if the fact is apparently contradicted (either by experience or another fact). Tracing back the origin or derivation of knowledge is the only way of assessing the cause and nature of mistakes.

By naturalising the account of truth so that it becomes the collective and verifiable counterpart to personal knowledge, we are able to examine some of the mechanisms by which it can be identified (i.e. learnt) in the first place. In this way we may be able to come to a fuller picture of truth that may be more useful because we can be more aware of the mechanisms which it relies upon.

9.   Recovering the utility of an "absolute framework"

Given all of what I have said why has the fiction of universal truths been so popular? I think the answer to this lies in the nature of philosophy. An absolute framework with universal truths simplifies discussion to an enormous degree if one can use the fictions of an absolute framework and absolute truths. It gives the impression that the content of the discussion has universal applicability – the idea is that somehow one could establish the universal truths first and "add in" the messy details later[7]. It is a defensible position against counter-examples, which is one of the driving dynamics of philosophical argument.

Thus I suspect the fiction of an absolute framework is rooted firmly in philosophical *practice*. It facilitates philosophy to a huge extent – in fact it is almost

---

[7] This is the implicit theory, in practice it almost never happens.

as if it forms the very playing field on which philosophy is done. It would be useful to see if some of this utility can be retained without the misleading downsides of the absolute framework.

If context can be characterised as an abstraction of the background/assumed conditions of a model (Edmonds, 1999), then this indicates that a truth can be useful in one of two circumstances: that the conditions under which it holds can be reliably recognised; or that these conditions are so pervasive that they can be safely assumed under "normal" conditions. In the second of these cases it can be helpful to cut out repeated reference to the normality conditions and simplify communication through the use of "philosophical context", which is the context *within* which all the events and discussion occur. This is a simulation of a universal framework because it is sufficiently wide to encompass all truths that normally hold. Such a context allows maximum generality and can be very useful in expressing objective truths in these conditions (sometimes called a "third person viewpoint").

What it can not do is *contain* counter-examples that go beyond the normal conditions it is situated in – the sort of counter-examples that philosophy abounds with (e.g. a landscape with facades of red barns obscuring real red barns, so that someone passing through believes, for good reason, that there is a red barn there, but is only correct by coincidence). This means that when these outer normality limits are reached and such a counter-example proposed, that the response is not to elaborate whatever proposition the example was made against, but to point the way in which the example was *abnormal*. One common way in which the normality conditions implicit in a discursive framework are breached in philosophy is via the means of reflection. Unlimited reflection can easily be incoherent and deceptive[8] in ways which can act to distort abstraction by placing unreasonable demands upon the (self) inclusiveness of expression[9].

10. The context of strong contextualism

A legitimate question to ask of the position espoused in this paper is "What is *its* context?" As one might expect, I think this is a highly legitimate (and probably pertinent) question of any argument that even hints at universality.

The original context in which the ideas were conceived was that of listening to philosophical discussions where I repeatedly observed how the insistence on universality (enforced by the means of abnormal counter-examples) resulted in conceptions that were inapplicable to the world. The ideas were further developed when I began to appreciate how knowledge might be developed via generalisation in a context when playing about with evolutionary learning algorithms. The context from which I have tried to write this paper is from a third-person viewpoint examining how we *actually* conceive of truth and how we might make our manipulation of truth more productive.

---

[8] See (Perlis, 1985, 1988) for some of the difficulties in unlimited self-reflection.

[9] An interesting example in philosophy where language and philosophy are discussed but without self-reference is (Wittgensein, 1975), which takes the safer and more certain option of the sort of "third party" context I am advocating.

11. Conclusion

I hope I have shown that a naturalised and context-dependent account of truth is not only credible and coherent but also allows for the development of a science of truth. In this way we can simultaneously reclaim the common sense usage of the word as well as deepening our understanding in it. The context of truths provides the "missing half" of the picture – the essential parts of truth that not explicitly represented, but are necessary for *meaning* and *the explanation of error*. The nature of truth can be ultimately traced to the contexts it was derived in (or the contexts from which it was abstracted from). Truth does not somehow *exist* in the abstract, separated from the messy contingencies of the world, but is very much part of the world.

In fact, all in all, the more pertinent question is not

"Why would one suppose that truth is context -dependent?"

but

"Why *on earth* would one suppose it wasn't *in the first place*?"

12. References

Adams, E. and Levine, H. P. (1975). On the Uncertainties Transmitted from Premises to Conclusions in Deductive Inferences, *Synthése*, **30**:429-460.

Barrow, J.D. and F.J. Tipler (1986) *The anthropic cosmological principle*. Oxford: Oxford University Press.

Cartwright, N. (1983). *How the Laws of Physics Lie*. Oxford: Oxford University Press.

Edmonds, B. (1999) The Pragmatic Roots of Context. CONTEXT'99, Trento, Italy, September 1999. *Lecture Notes in Artificial Intelligence*, **1688**:119-132.

Einstein, A. (1961), *Relativity: The Special and General Theory* translated by R. W. Lawson, New York: Crown Publishers Inc.

Giere, R. N. (1988). *Explaining Science: a cognitive approach*. Chicago : University of Chicago Press.

Glasersfeld, E. von (1995). *Radical constructivism : a way of knowing and learning*. London : Falmer Press.

McCarthy, J. and Buvac, S. (1998). Formalizing Context. (Expanded Notes). Unpublished manuscript. http://www-formal.stanford.edu/jmc/mccarthy-buvac-98/context/context.html

Penco, C. (1990). Objective and Subjective Context. CONTEXT'99, Trento, Italy, September 1999. *Lecture Notes in Artificial Intelligence*, **1688**:270-283.

Perlis, D. (1985). I: Foundations (Or: We Can Have Everything in First-Order Logic!), *Artificial Intelligence*, **25**:301-322.

Perlis, D. (1988). Languages with Self-Reference II: Knowledge, Belief and Modality, *Artificial Intelligence*, **34**:179-212.

Languages With Self-Reference I: Foundations (Or: We Can Have Everything in First-Order Logic!)

Artificial Intelligence, 25(3), pp. 301-322, 1985.

Quine, W. V. O. (1969). *Ontological Relativity and other Essays*, Columbia University Press, New York, 1969.

Weinberg, S. (1988). The first three minutes : a modern view of the origin of the universe. New York, NY : BasicBooks.

Wittgenstein, L. (1975). *On Certainty*. Oxford : Blackwell.

Young, R. (1999). Context and Supercontext. CONTEXT'99, Trento, Italy, September 1999. *Lecture Notes in Artificial Intelligence*, **1688**:417-441.