Restricting the Knowability Principle

Andrew Bacon

January 26, 2010

1 Introduction

Could there unknowable truths? Truths which, regardless of any extension to ones capacities or resources, remain impossible to know. The answer to this question is central in the evaluation of semantic anti-realism. But even for a metaphysical realist, the matter is far from closed.

Fitch's paradox has plagued the conscientious anti-realist, for whom every truth must at least be verifiable, since its introduction in the literature [6] as a problem for the knowability principle, below (henceforth KP.)

$$\forall p(p \to \Diamond Kp)^1 \tag{1}$$

According to to a theorem in Fitch [5], attributed to an anonymous referee[2], it follows that all truths are known:

$$\forall p(p \to Kp) \tag{2}$$

provided we assume a minimal logic of knowledge, namely

$$Kp \to p \tag{3}$$

$$K(p \wedge q) \to (Kp \wedge Kq) \tag{4}$$

i.e. knowledge is factive, and distributes over conjunctions. The proof goes by contradiction from assumptions (1), (3) and (4) and that for some q, that q is true, but unknown. See Figure 1 for the full derivation.

The conclusion, that all truths are known, is of course intolerable - even for an anti-realist. Accordingly, there have been many responses to the paradox. These mostly come in two varieties. Those which modify the background logic. Van Benthem [12] describes these approaches as like "turning down the volume on your radio so as not to hear the bad news". The other strategy is to weaken KP instead. Van Benthem likens these approaches to "censoring the news: you hear things loud and clear, but they may not be so interesting". In the first

[1] The use of propositional quantifiers can be eliminated in favour of a schema - it is presented as such for expositional convenience.

[2] The referee was later discovered to be Alonzo Church. See [REF].

F1. $q \wedge \neg Kq$ (Assumption)

F2. $(q \wedge \neg Kq) \rightarrow \Diamond K(q \wedge \neg Kq)$ (An instance of KP)

F3. $\Diamond K(q \wedge \neg Kq)$ (By F1, F2 and MP)

F4. $K(q \wedge \neg Kq)$ (assume for reductio)

F5. $Kq \wedge K\neg Kq$ (by (4) - knowledge distributes conjunction)

F6. $Kq \wedge \neg Kq$ (by (3) - factivity of knowledge)

F7. $\vdash \neg K(q \wedge \neg Kq)$ (by F4 - F7 reductio)

F8. $\Box\neg K(q \wedge \neg Kq)$ (by necessitation)

F9. $\neg\Diamond K(q \wedge \neg Kq)$ ($\Diamond$ is the dual of $\Box$)

F10. # (by F3 and F9)

Figure 1: A derivation of Fitch's paradox

category we have approaches which involve changing the logic to intuitionistic logic or paraconsistent logic (see Williamson [13] and Beall [1] respectively.) In the second camp there are restriction strategies such as restricting KP only to propositions of which it is consistent to suppose they are known (see Tennant [11]), "basic" propositions (Dummett - find reference), or actual truths (e.g. Edgington [4].) In this paper I shall be interested in a restriction strategy. The result is a modest knowability principle, however one that I believe is philosophically interesting, and well motivated from the perspective of a realist.

2 Transworld knowledge

Restricting KP clearly weakens the anti-realists position. However, since the anti-realist stance does not involve the thesis that all truths are known, this weakening is inevitable. The difficult problem is to identify a weakening that is properly motivated.

If the formulation of KP captured in (1) is supposed to reflect the anti-realist thesis that all truths are knowable, it had better comprise a correct analysis of knowability. In particular, it should turn out that a proposition is knowable just in case it is known in some possible world. Joe Salerno [REF] has already cast doubt on this equivalence, noting that expressions in English intended to express the knowability of a proposition are often factive. Being possibly known, on the other hand, is not; there is a world where I know I'm 6 meters tall because I am 6 meters tall in that world, and I've measured myself. But as it turns out I'm only 1.7 meters so I *can't know* that I'm 6m tall, because I'm not. That I am 6m tall is not knowable for me.

Although these cases are, on first looks, quite strong, I think it is worth being a bit careful that the seeming factivity is not just a pragmatic phenomenon. Salerno shows the insinuation of truth cannot simply be Gricean implicature, since it is not cancellable:

John could have known that he was being conned, even though no-one is in fact conning him. (5)

(5) is clearly marked. However it, it is still not clear that this is because truth is a *semantic* entailment, rather than simply a presupposition. Indeed, the insinuation of truth is preserved in negated and interrogative contexts, which strongly suggests we are dealing with presupposition

John couldn't have known that he was being conned. (6)

Could John have know he was being conned? (7)

I think there are independent reasons to be suspicious of the KP analysis of knowability. The counterexample I have in mind involves a proposition that *seems* to be knowable, which would be unknowable according to the KP analysis. Assume for the sake of argument that (8) is true:

The number of hairs on my head is even but nobody knows this. (8)

I claim that (8) is knowable. This is how I might go about finding out if (8) is true. Tomorrow I would count the hairs on my head. Assuming that you don't lose or grow hairs over night, I may infer that I had the same number of hairs yesterday and that noone could possibly have known this yesterday. Thus, I would know that I was even haired yesterday, and that noone knew I was even haired yesterday - and this constitutes knowledge of the proposition that, yesterday, I would have expressed by uttering (8).

On the other hand, according to the KP analysis of knowability, (8) is not knowable. It follows, given a minimal logic of knowledge, that (8) is necessarily unknown - for after all, it is a Fitch sentence and is susceptible to Church's argument in §1. The upshot of Fitch's paradox, in this case at least, seems to be that KP is an inadequate analysis of the knowability principle.

So we have established that (8), while intuitively knowable, cannot be known in any possible world for it would otherwise turn out that somebody *did* know the number of hairs on my head was even in that world. Yet the further conclusion that 8 is unknowable seems to be an artefact of our poor formalisation of knowability as 'known in some possible world'. We are concerned with the fact that I am even haired, and the fact that nobody knows this. However, in the worlds we are considering when evaluating KP , *this fact doesn't exist!* When we move to a different world, the facts change. Crucially, facts about what people know have changed - it has stopped being the case that nobody knows I am even haired. But what we were interested in knowing was the original fact in the actual world, not the new facts in a merely possible world. As Williamson [15] puts it

> The picture naturally associated with the claim that all truths are knowable is that if a truth about some subject matter is unknown, then that epistemic position could be different without any difference in the subject matter itself. [...] one might object, a difference in the epistemic position makes no difference to how the subject matter was *in the original world in which it was unknown*, as opposed to the world in which the epistemic position is different. In particular, if p is true in this actual world, why should it not be known in some non-actual world w that p is true in this actual world, rather than in w? – Williamson pp290-291

To take a specific example, consider knowledge of the proposition

The postmaster general invented bifocals. (9)

In this world, this knowledge is simply knowledge of the fact that Benjamin Franklin invented bifocals. However in other possible worlds this might be knowledge that someone else altogether, Joe Bloggs, invented bifocals, i.e. knowledge of a completely different fact. Edgington [4] suggests we overcome this problem by indexing the truth in question to the actual world. Since what obtains in the actual world does not change as we move from world to world, we can focus on a particular fact, from the perspective of a world w, without worrying about what obtains in w. With this in mind Edgington suggests the following amendment to KP (call this AKP.)

$$\forall p(@p \to \Diamond K@p) \tag{10}$$

On first sight this seems to be a large concession to Fitch. Of all the truths there are, we have restricted ourselves only to knowledge of propositions of the form @p, knowledge of necessary truths. However this is less of a compromise than it first looks, since any proposition p can be seen to be equivalent to @p *a priori*. This thought can be expressed using the $\Im$ operator (see [3])[3]:

$$\Im(@p \leftrightarrow p) \tag{11}$$

3 Problems with transworld knowledge

The approach adopted in this paper owes a lot to Edgington's [4] formulation of knowability in AKP. However there are several problems with AKP as it stands, the most pressing being its commitment to 'transworld knowledge' - knowledge that merely possible knowers have of actual truths. On the one hand, non actual knowers have to specify the actual world in such a way as is suitable for *de re* knowledge that p, say, is true in it. This proves problematic on a picture of worlds as causally isolated concreta, or causally inert abstracta. On the other hand, if we think of worlds as complete, informationally saturated scenarios, we must explain how they can be suitable to figure in the thoughts of knowers of finite epistemic means.

[3] $\Im p$ means, roughly, that p will be true no matter how the world actually turns out to be.

3.1 Specifying worlds

If we consider the the temporal analogue of the knowability paradox[4] and the analogue of Edgington's AKP[5] - it is relatively clear how a non-present knower could have knowledge of what is presently the case. Past situations leave traces for what is yet to come. Just as the bones of dinosaurs allow us to know what happened years ago, photos and so forth will allow future knowers to know what is presently the case. Similarly past beings might know what is presently the case by applying science to regularities in their past - much like we can predict the weather.

The picture is not quite so neat when we move to the modal case. Typically non actual worlds do not interact with the actual world in the way that past or future times can with the present. There are, however, a variety of ways to pick out the actual world which we survey below.

Indexicals. Clearly a non-actual knower cannot express their knowledge that actually p by simply thinking 'actually p', since since this will refer to the knower's non-actual world.

'Actually' as non-indexical. Suppose that we do not treat 'actually' as indexically picking out the world of utterance. Suppose an utterance with the same meaning as 'actually p' picks out the actual world in *every* world, rather than the world of utterance. Then de dicto knowledge of 'actually p' would be sufficient for transworld knowledge. However, this comes at a high price. There are many strange consequences of this view, for example, how do we know our use of 'actually' picks out *this* world as opposed to some other world? See Percival [9, pp92-93] for more discussion.

Demonstratives. Similarly non-actual knower's cannot simply indicate the actual world ostensively since they should be causally isolated (they should be parts of distinct possible worlds).

Names. Suppose we were to call the actual world Fred. Then you could say that knowledge of actually p can be expressed by saying 'p is true in Fred'. This similarly has difficulties, for example on a causal chain theory of reference, uses of the name Fred has to be related to the actual world though an appropriate causal chain. However this cannot happen since Fred is causally isolated from all non actual situations. At some point the actual world needs to have been picked out without the use of a name.

Definite Description. Another option would be to form a definite description which describes the actual world completely (one way to do this would be to conjoin all the atomic propositions true in the actual world.) So to have non-actual knowledge of the actual world the non-actual knower could say 'p is true in the world such that ϕ'. However, if ϕ does pick out the actual world uniquely and p is actually true then p should follow logically from ϕ. This seems to trivialise the type of knowledge required for transworld knowledge. Even if

[4]That every truth will be known at some time entails that every truth is known now. See Edgington [4].

[5]That for any truth that is now the case, it will at some time be known that it was now the case.

p follows from ϕ logically but non-trivially, there will be *another* description of the actual world in which it does so trivially, namely 'the world such that ϕ, and p holds'. If ϕ adequately picks out the right situation and p is true in that situation, then adding 'and p holds' should be superfluous. So this new description should pick out our world just as well as the original description.

Counterfactuals. Edgington's preferred solution is to pick out the world in question using counterfactuals. For example, suppose in a fit of insanity I decide to count every hair on my head. Call the world in which I do this w. After counting the hairs on my head, it occurs to me that the truth I have just uncovered is so exotic that if I hadn't counted the hairs on my head, no-one would ever have known they were even. In particular, if the actual world is the closest world to w in which I didn't count my head hairs, then I would have had knowledge that, in the actual world, I am even haired, but nobody will ever know this. As Williamson [15] describes it:

> Knowers in w need not describe x in detail. They describe one difference, q, from their world and let everything else be as close as possible to their world in the manner of counterfactual suppositions; they need not even know what their world is like. – Williamson, *Knowledge and its Limits*, 2000, p295

So, it might be thought that to have non-actual knowledge of the actual world is to have non-actual knowledge of the right sort of counterfactual (a counterfactual that singles out the actual world). However, Williamson [14] goes on to show that this condition is satisfied too easily. He shows that there will always be a counterfactual of this form but which is a trivial logically true counterfactual, assuming logical omniscience the trivial counterfactual for each world will be known in that world. Let's recap the proposal for counterfactual knowledge:

- It is known in world w that p is true in w* iff there is a q such that:
    1. K(q $\Box\!\!\rightarrow$ p) holds in w.
    2. w* is the closest world to w in which q is true.

Suppose that there is an r such that w* is the closest world to w in which r is true, and p is true in w* (that there is an appropriate r is a sufficiently weak assumption and that p is true in w* is mostly a matter of stipulation). So $(r \Box\!\!\rightarrow p)$ is true in w. But according to both Stalnaker and Lewis semantics for counterfactuals, w* is also the closest world to w in which $(r \wedge p)$ holds. Also assuming logical omniscience[6], $((r \wedge p) \Box\!\!\rightarrow p)$ is known in w since it is a logical truth. So if we let $q = (r \wedge p)$ we can see that both of the conditions above obtain. However, this has followed from the relatively weak assumption that there is a proposition r such that w* is the closest world to w in which r

[6]Even if we deny logical omniscience, we wouldn't want to say that having non-actual knowledge of actual truths is as simple as having knowledge of logical truths.

is true. So it follows that the counterfactual approach is excessively permissive, knowledge that p is true in w* will obtain in most worlds. Perhaps the problem is made clearest when we consider the case $w = w^*$. Any q that is true in w would do to satisfy 2. Then according to Edgington's condition, knowing that

> If I were even haired, then I would be even haired and nobody would know it. (12)

is sufficient for knowing 8. That is, knowing 9 is sufficient for knowing a Fitch sentence. But surely this cannot be the case: knowing a Fitch sentence leads to contradiction - this much we *can* take away from Fitch's proof (see F4-F7.) Since even Edgington accepts that Fitch sentences cannot be known in the actual world, this thesis cannot properly serve her purposes.

3.2 Completeness of worlds

There is a second problem for transworld knowledge. Worlds typically contain an infinite amount of detail. Unless we restrict ourselves to very trivial worlds, we need some kind of account of how we can have thoughts of the form *p holds in w.* If worlds are meant to be maximally informative, in that for any p, either *p holds in w* or $\neg p$ *holds in w*, then it is doubtful that worlds ever could figure in the thoughts of finite beings, let alone whether they ever do.

At this point Edgington, quite rightly in my opinion, suggests we should frame the theory in terms of knowing something holds in a *situation* rather than a world.

> When I think of the possibility that I will finish the paper today, I am not thinking of one totally specific possible world. It is not the sort of thing I am capable of thinking of. It, itself, seems to violate the principle of knowability. Nor am I thinking of a large class of possible worlds in which I finish the paper. I am thinking of a possibility or a possible situation, which I can refine, or subdivide, into more specific possible situations if I wish, but which will never reach total specificity. – Edgington, *The Paradox of Knowability*, 1985 [4]

So situations, unlike worlds, need not be complete in the sense that, in a world, every proposition or its negation obtains. Thus the situation Edgington describes above may feature in her desire to finish her paper in a way that no possible world could. However, the move to situation theory comes at a price. For one thing, we will have to drop AKP, for even if we interpret our modal logic within situation theory, the logic of actuality standardly entails the completeness of the actual situation[7]. So our first point of departure from Edgington will be to go overtly situation theoretic. Secondly, while the situation Edgington describes for us contains only a finite amount of information, there is no

[7] See axiom A2 in [2]: $(p \leftrightarrow \neg @ \neg p)$.

reason why some situations shouldn't be infinite. Indeed, we may think of possible worlds as maximal or complete situations, so we are back to square one if we allow a knowability principle that entails we can know infinitely complex statements such as *p holds in w.* Thus our next move is to propose a modest knowability principle, in which, if a proposition is true in a "suitable" situation, we should, in principle be able to know it is true in that situation. For after all, as realists interested in a knowability principle, why should we accept the full strength of KP , when all we are interested in is whether there is an epistemic concept which overlaps with truth in a compelling way?

4 Going situation theoretic

The idea, then, is to a develop modest principle of knowability that could be acceptable for a realist. To do this, we are going to have to go situation theoretic from the start, although we should be able to translate the general idea into a more traditional possible worlds framework. For now we shall not worry about the ontology of situations, whether they are sets of facts, mereological parts of worlds, events, location/situation type pairs, states of affairs and so forth. It is enough, for our purposes, that we have a set of individuals, Sit, to be thought of as our situations, however they are construed. Situations may also be more, or less specific than others by containing other situations as parts. For example my current situation, in which I am sitting at my laptop typing, is extended by the further situation in which it is also a miserable day outside. Sometimes one and the same situation can be extended in two incompatible ways, for example it could have been sunny outside while I am at my laptop typing, but at least one of these must be a non-actual situation. Regrettably, the second. So on top of Sit we can identify an ordering relation, $\sqsubseteq$, between elements of Sit. We postulate the following minimal restriction

Claim $\langle \mathsf{Sit}, \sqsubseteq \rangle$ is a partial order. (13)

Notice that in the example I just gave, we have a non-actual situation extending an actual situation. Namely my typing at my laptop (actual), while its being sunny outside (non-actual). This gives us a clue as to how to answer the first problem of transworld knowability. For example, there could be a non-actual knower (outside in the sunshine) peering in through my window into the actual situation - me at my laptop, typing[8]. The language just used seems very dramatic: are there really windows leading to non-actual worlds? Of course, me, and my creepy non-actual friend, are parts of two distinct possible worlds so it is not that one can simply move between them. Rather the assumption, when we adopt a situation theoretic approach, is that worlds overlap in a big way. On top of the occasional transworld identity, entire chunks of history can exist as parts of distinct worlds.

[8]This suggestion was also made in passing, by Lindström [8], although I have not seen it seriously developed anywhere since.

Enough of ordinary propositions, what of Fitch sentences, sentences of the form *p but nobody knows it*? Well, the reply is much like Edgington's. Suppose I am once again sitting indoors as I am now, and my perverse non-actual onlooker, who is getting creepier by the minute, has counted the hairs on my head. Now, he knows that in my situation I am even haired. Furthermore, he is aware of his exceptional creepiness and thus concludes that no-one in my situation would care to know this trivia. The fact that *he himself* knows it, is of no consequence for he is, thankfully, not a part of my actual situation. So a tentative knowability might go as follows, SitKP: if p is true in a situation s, then there is some situation *that extends* s in which it is known that p holds in s.

$$\forall s(In(s,p) \rightarrow \exists s^+ \sqsupseteq s(In(s^+, K(In(s,p))))) \tag{14}$$

The fact that s^+ extends s, allows us to avoid the pitfalls of AKP. Non-actual knowledge of actual facts can be as simple as knowledge gained by observation. This allows us to avoid transworld knowledge altogether since s will always be part of the same world as s^+. However, SitKP is not quite good enough. It is not suited to handle the second problem we have discussed: there is no restriction what kind of situations s can range over. If s takes values that are infinitely detailed (this notion will be made clearer later), then we must explain how $K(In(s,p))$ in the consequent of SitKP can obtain for finite knowers. But this is not the only problem we face if we let s range over every situation.

Definition 4.0.1. *Let* $w \in \mathsf{Sit}$. *We say that* w *is a* **maximal situation**, *or a* **possible world**, *iff* $\forall s(w \sqsubseteq s \rightarrow s = w)$.

Suppose we make a further reasonable restriction on structure of situations (refer to definition 5.0.3 p11):

$$\textbf{Claim } \langle \mathsf{Sit}, \sqsubseteq \rangle \text{ is a } \mathsf{cpo} \tag{15}$$

Theorem 4.1. *Possible worlds exist.*

Proof. By Zorn's Lemma. □

Theorem 4.2. *Assume* SitKP. *Let* w *be a possible world and* p *a proposition. It follows that: if In(w, p) then, In(w, K(In(w, p))). That is, if p is true in w, then p is known in w.*

Proof. Suppose that $In(w,p)$. By SitKP there is a situation $s^+ \sqsupseteq w$ such that $In(s^+, K(In(w,p)))$. By Definition 4.0.1 $s^+ = w$. Thus $In(w, K(In(w,p)))$. □

In particular, if we let w be the actual world, then by Theorem 4.2, it follows that all truths are known. We have the Fitch conclusion without even appealing to Fitch's proof! Notice, however, that this problem stems from the fact that in SitKP we allow s, the situations which we are supposed to figure in knowledge attributions, range over *all* situations, including possible worlds and other situations of infinite detail. If we were to restict ourselves to situations suitable for human knowledge, a notion we shall make precise later, this particular version of the paradox will be blocked.

5 Suitable situations

For a moment, let us a take a step back and ask what it is that a realist would want with a principle of knowability. It is certainly not to elucidate the notion of truth, since for a realist that notion is already tied to the world. Thus a knowability principle, if it were true, would provide a deep and interesting connection between knowledge and the world. It would be a substantive philosophical claim, not, as the anti-realist would have it, a conceptual truth. And as realists we should not shy away from rejecting the idea that reality stops short when we reach the contingent limitations of human thought. A moderate knowability principle, then, need only connect truth and knowability without further expecting it to limit the complexity of the world to the capacities of the human brain, or to insist the human mind can survey the full infinite detail of the world. If we wanted such a strong principle we would soon fall afoul of other problems. For example if we allowed the possible contents of our thoughts to be arbitrarily complicated and disjunctive we would be susceptible to Kaplan's paradox [7], which rules out the thinkability (and hence knowability) of arbitrary propositions through a Cantor style diagonalisation argument.

The idea, then, is this. To distinguish a certain class of situations that are, in some sense to be explained, 'finite' and suitable as subjects of knowledge. Then our knowability principle will say that whenever p is true in a 'graspable' situation, s, we can come to know that p is true in that situation. Moreover, there will be a situation that extends s in which it is known that p is true in s. To be sure, this principle is timid compared KP , but it is a substantial claim nonetheless, and to my mind, one that is neither obviously true, nor obviously false.

To arrive at our notion of 'graspable' in a general way, i.e. a way that doesn't assume a particular ontology of situations, we must make a small detour and introduce some order theoretic definitions.

Definition 5.0.1. *Let $\langle P, \leq \rangle$ be a partial order. If $S \subseteq P$ then we say s^* is an* **upperbound** *for S iff $s \leq s^*$ for each $s \in S$. We say s^+ is a* **supremum** *for S iff s^+ is an upperbound for S, and whenever s^* is an upperbound for S, $s^+ \leq s^*$. We write $\bigsqcup S$ for the supremum of S, if it exists.*

Definition 5.0.2. *Let $\langle P, \leq \rangle$ be a partial order and let $S \subseteq P$. S is* **directed** *iff whenever $s, s' \in S$, $s \leq t$ and $s' \leq t$ for some $t \in S$.*

Definition 5.0.3. *A* **cpo***, $\langle P, \leq \rangle$, is a partial order with the property that whenever S is a directed subset of P, S has a supremum in P.*

Definition 5.0.4. *A element $c \in P$ is* **compact** *iff for every directed set S, if $c \leq \bigsqcup S$, then $c \leq s$ for some $s \in S$.*

With this general machinery behind us we are in a position to apply this to our case of interest: $\langle \mathsf{Sit}, \sqsubseteq \rangle$. First we must make the following important assumption about $\langle \mathsf{Sit}, \sqsubseteq \rangle$[9]

$$\textbf{Claim } \langle \mathsf{Sit}, \sqsubseteq \rangle \text{ is a } \mathtt{cpo} \tag{16}$$

The key concept here is that of a *compact situation.* Compact situations are supposed to be those situations that are graspable and knowledge apt.

Definition 5.0.5. *A situation c is* **compact** *iff for every directed set S of situations, if S has a supremum, $\bigsqcup S$, and $c \sqsubseteq \bigsqcup S$, then $c \sqsubseteq s$ for some $s \in S$*

To see how this cashes out for a particular way of construing $\langle \mathsf{Sit}, \sqsubseteq \rangle$ let's interpret situations as consistent sets of formulas in some first order language, and $\sqsubseteq$ as simply the subset relation, $\subseteq$. Situations are almost certainly more complex than this, but this will help us fix ideas. Note the compactness theorem for FOL guarantees this forms a $\mathtt{cpo}$. The following lemma confirms that the notion of compact situation is the one we are looking for in this case:

Lemma 5.1. *Let s be a situation. The following are equivalent*

- *s is compact*
- *s supports finitely many formulae*

Proof. Suppose s is a finite set of formulae, and suppose $s \sqsubseteq \bigsqcup S$ for some set, S, of situations. Then whenever $\phi \in s$ there is a $t_\phi \in S$ such that $\phi \in t_\phi$, since $s \subseteq \bigcup S$. s is finite, so $\{t_\phi \mid \phi \in s\}$ is finite. Since S is directed there is a $t \in S$ for which $t_\phi \sqsubseteq t$ for each $\phi \in s$. If $\phi \in s$ then $\phi \in t_\phi \subseteq t$ - therefore $s \sqsubseteq t \in S$. Therefore s is compact.

Conversely, let s be compact. Let S be the set of finite subsets of s. The members of S are consistent since s is. S is directed since if s_1 and s_2 are finite subsets of s so is $s_1 \cup s_2$. Clearly S has a supremum, $\bigsqcup S = s$. $s \sqsubseteq s = \bigsqcup S$ thus $s \sqsubseteq t$ for some $t \in S$ by compactness of s. $t \in S$ so t is finite, thus s must be finite too. □

Corollary 5.2. *There are no compact maximal situations. Moreover, there are no maximal compact situations.*

The significance of this last result is that it blocks the conclusion of Theorem 4.2. For every compact situation, there is a situation *strictly bigger*, and so there is always room for more facts to be known of them. We are finally in a position to give a refined version of SitKP. Let C be the set of compact situations.

$$\forall s \in C(In(s,p) \rightarrow \exists s^+ \sqsupseteq s(In(s^+, K(In(s,p))))) \tag{17}$$

[9] This claim is not incontestable, see for example Menzel (find reference.) His counterexample applies to states of affairs, and it appears as if it can be avoided with a fleshed out semantics for $In(s, Qx\phi)$ where Q is a generalised quantifier. For example $In(s, \exists x\phi)$ is true iff for some a, $In(s, \phi(a))$ is true iff for some a, $a \in [\![\phi]\!]^+_s$. Similarly $In(s, \exists x\neg\phi)$ is true iff for some a, $a \in [\![\phi]\!]^-_s$ - the anti-extension of ϕ at s.

6 Similar views

Lindström [8] develops a situation theoretic approach to Fitch's paradox much like the one presented here. When it comes to transworld knowledge Lindström, suggests, as on this account, that knowledge across situations is possible simply through perception. However, when faced with the second Fitch paradox (Theorem 4.2) Lindström recommends we take situations to be always 'partial and possible to extend', as a matter of their metaphysics. As such there would be no possible worlds, contradicting Theorem 4.1[10]. Aside from the technical difficulties this would face, there is also considerable pressure from elsewhere in philosophy and linguistics to accept possible worlds.

Rabinowicz and Segerberg [10] defend Edgington's AKP from within a traditional possible worlds framework. The problem they address is this: the fact that if p is true, $@p$ is true in all worlds, and on a standard Hintikka semantics for K, entails that $(p \rightarrow K@p)$ for all p. Fitch strikes again. Lindström [8] suggests that this problem stems from the fact that in epistemic logic we assume the agents are perfect logicians. But if p is an empirical fact, it is hard to see how a logically perfect agent could come to know @p from her armchair if she didn't already know p. Rabinowicz and Segerberg take a more promising line and resolve the issue in a two dimensional semantics. This general strategy would be useful in the present account. Indeed the idea is much in line with a lot of work that has already gone on in situation semantics, so we won't attempt to clutter the literature here. However even this might not be necessary. It certainly seems that SitKP deals only with necessary truths much like AKP - a supposed defect in the theory. However, if we think of situations as events, or some other type of individual, there is no need to insist they have their properties necessarily. For example it might only be contingently the case that $In(s, p)$. On first glance, this suggests that we can't understand contingency in terms of truth in some situation. So one way to go would be to develop an additional theory of worlds, in which situations figure as individuals like you or me. But perhaps we needn't even do that. We might say $In(s, p)$ is possible iff for some situation s', $In(s', In(s, p))$.[11] Unfortunately, a detailed account along these lines would take us too far afield to present here.

7 Conclusion

We have discussed a modest knowability principle which does not suffer from Fitch's paradox. The principle is, of course, weaker than the full KP since it is

[10] There are still ways out at this point. You could deny that Sit forms a cpo. Or perhaps you could reject the implicit assumption that Sit is set sized and claim that situations are indefinitely extensible yet still closed under the suprema of directed sets, like the ordinals are. Lastly, you could deny the axiom of choice on intuitionist grounds - Williamson [13] provides further motivation for going intuitionist about Fitch's paradox.

[11] We would expect s to be in the domain of s' and so it would be plausible that this happens only when $s \sqsubseteq s'$. Thus, on a theory like this we would have to reject the principle that $s \sqsubseteq s'$ entails $\forall p(In(s', p) \rightarrow In(s, p))$.

limited to truth about small portions of the world - but a limitation, however, that is externally motivated. After all, it would not be surprising if an agent, who could all at once observe the world in its entirety, knew everything about it. Once we do away with such commitments the restriction to compact situations in SitKP is entirely reasonable. This being said, SitKP postulates a substantial connection between knowledge and the world and has the virtue of being neither obviously true, nor obviously false. For example, while it is conceivable that an observer could know I am even haired - even though it would require very large (but finite) extension to our ordinary resources. But what about knowledge of deep metaphysical questions? For example knowing whether time is continuous or merely dense? SitKP might even be contingently true: if time were discreet presumably this truth would be knowable. We could observe that no event endures longer than a chronon. But if it is dense you could argue this would be unknowable - for every event we observe of duration t, we have not ruled out that the smallest unit of time is $t' < t$.

With all this said and done, we haven't even touched on mathematical truths. Perhaps it is odd to think of the theorems of ZFC, say, holding in a physical situation. At least, if a theorem holds in any situation it holds in all of them. Thus it seems as if the restrictions we have discussed are superfluous on the case of mathematics. But the question of whether all mathematical truths are in principle knowable is a delicate one, as witnessed by Gödel's incompleteness theorems on the one hand, and the possibility of supertasks, fair nondeterminism, and analogue computers on the other.

References

[1] J.C. Beall. Fitch's proof, verificationism, and the knower paradox. *Australasian Journal of Philosophy*, 78:241–247, 2000.

[2] J Crossley and L Humberstone. The logic of "actually". *Reports on Mathematical Logic*, 8:11–29, 1977.

[3] Martin Davies and Lloyd Humberstone. Two notions of necessity. *Philosophical Studies*, 38(1), 1980.

[4] Dorothy Edgington. The paradox of knowability. *Mind*, 94(557-568), 1985.

[5] F Fitch. A logical analysis of some value concepts. *The Journal of Symbolic Logic*, 28:135–142, 1963.

[6] W. D. Hart and C. McGinn. Knowledge and necessity. *Journal of Philosophical Logic*, 5:205–208, 1976.

[7] David Kaplan. A problem in possible world semantics. In Diana Raffman Walter Sinnott-Armstrong and Nicholas Asher, editors, *Modality, Morality and Belief: Essays in Honor of Ruth Barcan Marcus*, pages 41–52. Cambridge University Press, Cambridge, 1995.

[8] Sten Lindström. Situations, truth and knowability: A situation-theoretic analysis of a paradox of fitch. In E. Ejerthed and S. Lindstrm, editors, *Logic, Action and Cognition: Essays in Philosophical Logic*, pages 183–210. Kluwer Academic Publishers, 1997.

[9] Philip Percival. Knowability, actuality and the metaphysics of context-dependence. *Australasian Journal of Philosophy*, 69:82–97, 1991.

[10] W. Rabinowicz and K Segerberg. Actual truth, possible knowledge. *Topoi*, 13:101–115, 1994.

[11] N. Tennant. *The Taming of the True*, chapter Chapter 8. Oxford, 1997.

[12] J. van Benthem. What one may come to know. *Analysis*, 64(2), 2004.

[13] T. Williamson. Intuitionism disproved? *Analysis*, 42:203–207, 1982.

[14] Timothy Williamson. On the paradox of knowability. *Mind*, 96:256–61, 1987.

[15] Timothy Williamson. *Knowledge and its Limits*, chapter 12. Oxford, 2000.