Semantic Paradox and Alethic Undecidability

Steve Barker, University of Nottingham: stephen.barker@nottingham.ac.uk

In what follows, I use the principle of truth-maker maximalism, *TM* below, to provide a new solution to the semantic paradoxes:

TM: If a sentence is true (or false), then it is true (or false) in virtue of non-alethic facts.

In *TM*, the term *non-alethic facts* means facts that have nothing to do with the truth or falsity of sentences, like the fact that snow is white. The solution I propose works this way. Call *being grounded* the property a sentence *S* has when it has a truth- or falsity-maker that is a non-alethic fact, and *ungrounded* the property *S* has when it lacks such a truth- or falsity-maker. A common view about paradoxical sentences, for example, *L*, is that they are *ungrounded* and thus, given *TM*, that they lack truth-value:

L : *L* is not true

I show below that the argument concluding that *L* is ungrounded is flawed. I demonstrate that we can neither *affirm* that paradoxical sentences are grounded, nor *affirm* that they are ungrounded. It's undecidable whether they are grounded or not. That's due to their grounding status being *infinitely deferred* as a result of the looping or infinitely descending referential chains these sentences exhibit. (I say more about those below.) Given *TM*, it follows that paradoxical sentences are *alethically undecidable*, that is, we cannot assert, in principle, whether they are true, false, either true or false, neither true nor false, both true and false, and so on. Call this the alethic undecidability solution or *AUS*.

AUS doesn't face the immediate threat of *revenge* that a truth-value gap approach encounters. The latter claims that *L* is neither true nor false, and so, *L* is not true. But since *L* says of itself that it is not true, it must be true after all, contradicting the initial claim of the gap-approach. Since, given *AUS*, we refrain from assigning any truth-value to *L*, including *neither-true-nor-false*, we are not committed to *L*'s lacking truth, and so escape this problem.

No other obvious forms of revenge threaten *AUS* because of its stance that paradoxical sentences are undecidable with respect to all alethic properties.

Below, I clarify my core claims and then argue for them.

1. Let an *alethic predicate* be any one-place open-sentence of the form *(...x ...true...)* or *(...x ...false...)* where the other constituents of the open-sentence are standard logical operators: '&', '¬', '∀x', etc, or standard modal operators. So, alethic predicates include, *x is true*, *x is false*, (¬¬x is true), and so on. Alethic predicates do not include intentional verbs like *believe* or *assert*. So, *(believes of x that x is true)*, *(x is asserted to be true)*, or *(x cannot be said to be true)* are not alethic predicates. Call an *alethic property* any property corresponding to an alethic predicate. We shall be relaxed about the metaphysics of properties in what follows and allow that each meaningful predicate corresponds to a property. (Nothing hangs on it.)

By *alethic fact* I mean any fact about a sentence or sentences that can be described using only alethic predicates and reference to sentences. *TM* above distils the idea that the truth or falsity of sentences, whether contingent or necessary, is ultimately determined by *non-alethic fact*, in this sense. I will not be concerned with the general metaphysical question of how exactly to fashion an adequate development of *TM* or with defending the principle.[1] *TM* is controversial but many find it compelling.

Paradoxical sentences, sentences like *L*, are members of a class of sentences with looping or infinitely descending referential chains such that it is impossible to derive any specification of a ground for such sentences. Below, *Regress* and *Loop,* and *T* are other instances of this class:

Regress:	(0) (1) is true	*Loop*:	(0) (1) is true
	(1) (2) is true		(1) (0) is true
	(2) (3) is true		
	etc.	*T* :	*T* is true

---

[1] *TM* itself must, by its own lights, have a truth-maker. The truth-maker of *TM*, we may suppose, is the very nature of truth/falsity to be dependent in the way indicated by *TM*. This is a fact about truth, but not an alethic fact in the sense we have defined. I won't examine here how we ought to conceive of the nature of truth/falsity.

Regress is a clear instance of the unspecifiability of ground I am pointing to. In *Regress*, sentence (0) says *(1) is true*. (1) says *(2) is true*. (2) says *(3) is true*, and so on. Each sentence in the list simply says that the sentence below it is true. At no point do we find a sentence that describes a non-alethic condition—like *snow is white*, or *grass is purple*. Similar comments apply to *L* and *T*, the only difference being that *L* and *T* are self-referential. An inquiry into what *L* is ultimately about always loops back to *L*, a sentence describing the lack of truth of a sentence. We never encounter a sentence that specifies a ground, like snow is white, since we always return to *L*. *Loop* (0) is just like *T* and *L* in this respect, except there is an intermediate sentence.

Call the property exhibited by these sentences *being ground-unspecifiable*. So a sentence *S* is ground-unspecifiable if and only if there is no sentence, *referentially downstream* from *S* that specifies a potential ground for *S*.[2] Analytically the property of being ground-unspecifiable is distinct from the property of being *ungrounded*. The former implies that it is impossible to specify a ground for a sentence *S*. It does not follow, logically, from this, that there is no ground for *S*. It could be that the impossibility of specifying a ground for *S* arises not from the fact that *S* lacks a ground, but from the fact that the grounding status of *S* is *undecidable*. That would mean that even a perfect representation of reality would neither represent *S* as grounded nor represent *S* as ungrounded. This, indeed, is what I argue is the case. Ground-unspecifiable sentences are undecidable as to whether they have grounds or not. We cannot then affirm that they are ungrounded.

In what follows I first present the case for ground-unspecifiable sentences being ungrounded. I will then show how the case is flawed, and demonstrate why we have to accept undecidability.

2. It seems plausible to think along the following lines. The only way a sentence *S* can be grounded is if there is a sentence *R* referentially downstream from *S* that describes a non-alethic condition. (That's how *S* gets hooked onto non-alethic reality.) Contraposing then we ought to accept:

---

[2] I won't define 'referentially downstream' here, since the intuitive idea is obvious enough.

G-principle: If there is no sentence *R* downstream from *S* that describes a non-alethic condition, then *S* is ungrounded.

Ground-unspecifiable sentences are such that no sentence *R* referentially downstream from them describes a non-alethic condition. (For example, inspect *Regress* (0). All downstream sentences describe alethic conditions.) So, given *G-principle*, ground-unspecifiable sentences are ungrounded.

Here is another closely related argument for ungroundedness. Think of a hierarchy of sentences modelled on Kripke's (1975) hierarchy. Level-0 sentences are about non-alethic reality. Level-1 sentences are truth- or falsity-ascriptions of whatever logical complexity about those sentences. Level-2 sentences are truth- and falsity-ascriptions of whatever logical complexity about level-1 sentences or about both level-1 and level-0 sentences, and so on, up to transfinite levels. If a sentence is in the hierarchy, it's grounded. But one might also hold:

H-principle: If a sentence *S* is not in the hierarchy, it's ungrounded.

Ground-unspecifiable sentences are not in the hierarchy, and so, by *H-principle*, are ungrounded. Accepting *H-hierarchy* boils down to accepting *G-principle*. The reason is that being in the hierarchy and having a sentence *R* referentially downstream describing non-alethic reality are equivalent. (That's easily confirmed by the reader.) Hence this argument adds nothing essentially new to the case for ungroundedness.

Should we accept these arguments? No. They are flawed. Affirming *G-principle* and *H-principle* only looks plausible if we ignore the infinite deferment of grounding that the dependency of truth and ground-unspecifiability imply. Let's first see what this infinite deferment is and then how it impacts on our acceptance of these principles.

Take (0) in *Regress*. In considering the question of whether (0) is grounded or not, one should be aware of the following facts about dependency relations between sentences in *Regress*. Suppose a sentence (*n*) in *Regress* is grounded. Then (*n* + 1) must be grounded, and, moreover, (*n*) is grounded *because* (*n* + 1) is, and not vice versa. Suppose (*n*) is ungrounded.

Then (n) must be ungrounded *because* ($n + 1$) is ungrounded, and not vice versa. Summing up we can say:

> For all sentences (n) in *Regress*, the grounding status of (n)—its being grounded or its being ungrounded—holds in virtue of the grounding status of ($n + 1$)—its being grounded or its being ungrounded.

This general fact of dependency just arises from the fact that every sentence (n) in *Regress* inherits its alethic features from ($n + 1$), since (n) is about ($n + 1$)'s truth.

Given these dependency relations, (0)'s being grounded/ungrounded depends on (1)'s being grounded/ungrounded, which depends on (2)'s being grounded/ungrounded, and so on, ad infinitum. Or, for all (n), (0)'s being grounded or not depends on (n)'s being grounded or not. What should we say about (0)'s grounding status? I submit that we should say that (0)'s grounding status is infinitely deferred. It's never settled, and so undecidable. We cannot then affirm, following *G-Principle*, that (0) is ungrounded. Let's see why exactly.

Consider (0) in *Regress*. It is true that no sentence (n) in *Regress* referentially downstream from (0) describes a non-alethic condition. But to infer—as *G-principle* requires—that (0) is ungrounded from that fact is to forget the infinite deferment of grounding. Given the referential structure of *Regress*, the grounding of no sentence (n) in *Regress* depends on that sentence describing non-alethic reality, but always on the grounding of the sentence ($n + 1$) below it. So no sentence (n)'s failure to describe non-alethic reality can be a reason for claiming that (0) fails to be grounded, since each sentence (n) transfers responsibility for fixation of grounding onto the next sentence ($n + 1$). Since no sentence (n) can be held up as the source of (0)'s failing to be grounded, we cannot say that (0) is ungrounded on the basis that no sentence (n) describes a non-alethic condition, which is what *G-principle* assumes.

In sum: All sentences (n) in *Regress* transfer the fixation of (0)'s grounding status to ($n + 1$). So, the fact that no (n) describes a non-alethic condition, cannot be a reason to affirm (0)'s ungroundedness since its grounding never depends on any of them doing so. Hence we cannot affirm that (0) is ungrounded. We cannot affirm that (0) is grounded either—that

would require uncovering a sentence *R* referentially downstream from (0) describing a non-alethic-condition. So (0)'s status is undecidable.

It follows that *G-principle* cannot be affirmed. Indeed, it's undecidable since there are sentences *S*, with no sentences *R* downstream from them describing non-alethic conditions, whose grounding status is undecidable. Since *H-principle* is equivalent to *G-principle*, the former must undecidable also. The Kripke-style argument then fails.

We have been looking at sentences like *Regress* (0) whose ground-unspecifiability arises from infinitely descending referential chains. But, we can apply the infinite deferment argument for undecidability to sentences whose ground-unspecifiability arises through looping, the simplest form of which is *T*. Implicit in *T*'s self-reference is an infinite series of interpretative stages. We can partially express this as follows:

> *T* refers to a sentence that says *T is true*, which refers to a sentence that says *T is true*, which refers to a sentence that says *T is true*, ad infinitum.

At each stage we ascertain a content: *T is true*. This content describes an alethic condition, that a sentence *T* is true. So, at each stage, no grounding condition is specified. But there is always a further interpretative stage after any given stage. This shows that *T*, due to its self–reference, is referentially down stream from itself.

Using *G-principle* and *H-principle*, we might attempt, again to show that *T* must be ungrounded. But again, for exactly the same reasons we can show these arguments are flawed. The fact that there is no sentence downstream from *T* describing a non-alethic condition cannot be a reason for claiming that *T* fails to be grounded, since *T* as each interpretive stage transfers responsibility for fixation of grounding onto *T* at the next stage. Since *T* at no stage can be held up as the source of *T*'s failing to be grounded, we cannot say that *T* is ungrounded on the basis that no sentence at any interpretive stage describes a non-alethic condition.

In short, whether sentence *S*'s ground-unspecifiability comes through infinite descent or looping we can argue for the same conclusion: it's undecidable whether *S* is grounded or not. *Undecidability* here just means unassertable in principle. This unassertability is not due

to epistemic limitation or verification transcendent fact. It means that a perfect representation of reality would neither assign such sentences the status *grounded* nor *ungrounded*.

3. Ground-unspecifiable sentences are undecidable with respect to their grounding status. Given *TM*—which says that all alethic properties must be grounded—it follows that such sentences are undecidable with respect to their alethic properties. For example, it's undecidable whether (0) in *Regress* is either true or false, lacks truth and falsity, is both true and false, and so on. Each alethic property, given *TM*, requires a non-alethic truth- or false-maker. But none can be assigned to (0). Call this undecidability with respect to alethic properties *alethic-undecidability* or ^{A}*undecidability* for short.

As we noted, the class of ground-unspecifiable sentences includes those sentences that are the alethic paradoxes, the simplest, and best known of which, are the Liars, sentences such as *L* above. So, given *TM*, it's undecidable whether *L* is true, false, either true or false, neither true nor false. In particular, it's undecidable whether *L* is both true or false. It's not denied, but it's not affirmed either. To be both true and false, by *TM*, *L* would have to have both a truth-maker and a falsity-maker, but it's undecidable that it has either.

That's the essence of the alethic undecidability solution—*AUS*.[3] *AUS* uses *TM* to derive the conclusion that *L* is ^{A}undecidable. What, however, of the derivation of a contradiction from *L* via the *T*-schema? To so derive a contradiction, we need to be able to affirm the biconditional T^L below:

T^L : *L* is true iff *L* is not true

However, T^L is ^{A}undecidable, since it's a ground-unspecifiable sentence, given *L* is. So, *AUS* undercuts the standard reasoning generating the contradiction.

Although the biconditional T^L is ^{A}undecidable, we can still affirm the validity of the *T*-schema. The reason why is that our judgements of the validity of logical/semantic schemata should not take into account ^{A}undecidable sentences, since ^{A}undecidability is not

---

[3] *AUS* is not a meaningless solution. We cannot affirm that *L* expresses no proposition, as some theorists have proposed (Goldstein 2000). If *L* expresses no proposition then it lacks truth. But we cannot say *L* lacks truth according to *AUS*. Therefore we cannot say *L* fails to express a proposition.

an alternative truth-value, but, instead, undecidability with respect to all truth-values. There is no *logical* incompatibility between *S is ^{A}undecidable* and *S is true* or *S is not true*, etc. Affirming, *S is true but ^{A}undecidable* is a form of Moore's paradox. It's like asserting: *S is true but one cannot assert that S is true*. We cannot therefore define logical connectives with ^{A}undecidability as one value. If ^{A}undecidability does not enter into logical relations with truth-values, it ought not to be considered in connection with validity. Validity should be defined purely in terms of ^{A}*decidable* sentences. In those terms, the *T*-schema is valid. Therefore, in refraining from asserting T^L, we are not tacitly withdrawing from the *T*-schema.[4]

The impact of ^{A}undecidability on reason and logic is purely *pragmatic*. If we discover that a sentence *S* is ^{A}undecidable—due to its being ground-unspecifiable—we ought to stop reasoning with it, unless, that is, we explicitly or implicitly assume that it is ^{A}*decidable*. If I assume *L* is ^{A}decidable, I can reason with *L*. But when I uncover contradictions, and recognize that it's ground-unspecifiable, I should just conclude that it's ^{A}undecidable, and withdraw from the reasoning.

4. AUS applies uniformly to paradoxical sentences always issuing in the judgement that they are ^{A}undecidable. Consider the contingent liars involving universal quantification, below, where on page 1, P^1 is written, and page 2, P^2 is written:

> P^1: Every sentence on page 2 is not true
>
> P^2: Every sentence on page 1 is true.

It's easy to show that P^1 and P^2 are ground-unspecifiable. For example, P^1 says that certain sentences on page 2 are not true. The only sentence on page 2 is P^2 and it says that the sentences on page 1 are true. The only sentence there is P^1 and it says the sentences on page 2 are not true. And so on. At no point can a specification of a condition be derived that does not concern the truth or lack of truth of sentences. Given the circumstances of the case, both

---

[4] Note that the fact that there are ^{A}undecidable sentences does not undermine our affirmation of *TM*. If it's *decided* that any sentence is true or false, then it's *decided* that it has a ground. If the alethic status of any sentences is undecidable, then it's grounding status is undecidable, and vice versa.

sentences are ground-unspecifiable. Our conclusion should not be that they are ungrounded and, and so lack truth-value, but that it's undecidable what the alethic properties of these sentences are.[5]

Consider now Curry's paradox:

C :   If C is true, then God exists

Given a supposition that C is true, we can easily prove that God exists, and so, by conditional introduction, prove C, and thus that C is true, and thus that God exists by modus ponens. How does AUS deal with this? The sentence C *is true* is a ground-unspecifiable sentence. It's ground-unspecifiable because C is. C is ground-unspecifiable because of its reflexive referential properties. As noted above, we should not allow proofs with ground-unspecifiable sentences. So, the right response to the Curry is just to refrain from the initial step of the proof, which is to suppose that C is true. The proof then is blocked.

AUS also explains what's going on in open pair sentences such as:

O_1 : O_2 is not true.
O_2 : O_1 is not true.

We can assign O_1 truth and O_2 falsity (or neither-true-nor-false), or O_1 falsity (or neither-true-nor-false) and O_2 truth, but we cannot assign the pair identical truth-values. The problem is that the sentences are semantically identical, and so should get the same truth-values.[6] In contrast to approaches that assign truth-values, AUS has no difficulties at all with open pairs. It gives a perfectly symmetrical and principled analysis. Each member of the pair is [A]undecidable for exactly the same reasons. Each is ground-unspecifiable. So, AUS gives us symmetry and consistency.

5. It seems that AUS might have a general solution on offer to the semantic paradoxes. However, every solution to the semantic paradoxes must confront the challenge of revenge.

---

[5] It's simple to see how AUS applies to Yablo's (1993) paradox. An infinite sequence of sentences each one saying that all below it are untrue will be such that all sentences are ground-unspecifiable, and so [A]undecidable. No paradox can then arise.

[6] See Armour-Garb and Woodbridge 2006 for detailed examination of open pairs.

The obvious sorts of candidates for revenge for *AUS* are sentences like sentence (0), (1), etc, in *I-Regress* and *R* in *I-Loop*:

I-Regress :   (0) (1) is ^{A}undecidable   *I-Loop* :   *R* :   *R* is ^{A}undecidable.

(1) (2) is ^{A}undecidable

(2) (3) is ^{A}undecidable

etc.

Are these sentences ^{A}undecidable? If they are, then it seems that *AUS* is in trouble. Suppose (0), (1), etc, are ^{A}undecidable. Then (0) says something of (1) that's true, namely, that it's ^{A}undecidable. So we conclude: *(0) is true but one cannot assert that it's true*. Perhaps that's not a formal contradiction, but at the very least it's a Moorean pragmatic one.

Fortunately, this problem never arises, since, in fact, the claim that these sentences are ^{A}undecidable is ^{A}undecidable. This is how. Ask whether (0) is ^{A}undecidable. Answers depend on the status of (0), and what it says. Its subject-matter is (1). To determine if (1) is ^{A}undecidable, we have to consider its subject matter, (2). And so on, ad infinitum. In short, determination of the outcome of the question *Is (0) ^{A}undecidable?* depends on the outcome of *Is (1) ^{A}undecidable?* which depends on the outcome of *Is (2) ^{A}undecidable?*, and so on ad infinitum. We enter an infinitely descending *processing regress* with respect to the initial enquiry. If this is right, it's undecidable whether the condition *(0) is ^{A}undecidable* obtains. That means one cannot say that it's true that one cannot say that (0) is true. And so on. So there is no threat of contradiction after all. (The same holds for *R-Loop*, except that in this case we have processing loop rather than infinite descent.)

One might object that even if it's ^{A}undecidable that (0) is ^{A}undecidable, we cannot avoid recognizing that we cannot affirm: *(0) is true*, *(0) is false*, *(0) is neither true nor false*, etc. Does that not mean we have to recognize that (0) is ^{A}undecidable? But noting that we won't in fact affirm, *(0) is true*, *(0) is false*, *(0) is neither true nor false*, etc, because it's ^{A}undecidable that (0) is ^{A}undecidable, is not the same as recognizing that we cannot affirm it. The fact that we do not affirm these sentences is the result of the fact that we cannot affirm

that we *can* affirm them, and we cannot affirm that we *cannot* affirm them. There is no affirmation that we cannot affirm them as such.

     AUS escapes these forms of revenge. Exploring whether it escapes others goes beyond the bounds of this paper.[7] My goal has merely been to show that given a highly intuitive principle about truth, *TM*, a novel and promising solution to the semantic paradoxes immediately follows without ad hoc modification of logic or obvious revenge.[8]

References

Barker, S. 2012. 'Truth-Making and the Alethic Undecidability of the Liar'. *Discusiones Filosoficas* 14: 21.

Armour-Garb, B and J. Woodbridge. 2006. 'Dialetheism, Semantic Pathology, and the Open Pair', *Australasian Journal of Philosophy* 84: 395 – 416.

Goldstein, L. 2000, 'A Unified Solution to Some Paradoxes,' *Proceedings of the Aristotelian Society*, 100, pp53-74.

Kripke, S. 1975, 'Outline of a theory of truth', *The Journal of Philosophy* 72: 53-81.

Yablo, S, 1993: 'Paradox without Self-reference', *Analysis*, 53, pp. 251-2.

---

[7] I provide an independent argument for *AUS* using the concept of truth-conditions and truth-making in Barker 2012.