Lecture 4: Is that Really Revising Logic?

König's paradox (and Berry's variant).

Let L be any language whose formulas are finite strings of finitely many basic symbols. Then

(K1) There are only countably many formulas of the language;

(B1) There are only finitely many formulas of the language of length less than 1000.

Say that object o is *definable in* language L iff there is a 1-place formula of L that is true of o and not true of anything else.

Say that object o is *1000-definable in* language L iff there is a 1-place formula of L *of length less than 1000 symbols* that is true of o and not true of anything else.

Each formula of L defines at most one object, so

(K2) There are only countably many objects definable in L

(B2) There are only finitely many that are 1000-definable in L.

But there are uncountably many ordinal numbers, and so

(K3) Given any L, there are ordinal numbers that aren't definable in L.

Similarly, there are infinitely many natural numbers, so

(B3) Given any L, there are natural numbers that aren't 1000-definable in L.

(K3) and (B3) are the bases for König's and Berry's paradoxes.

König's paradox: From

(K3) There are ordinals that aren't definable in L

infer (using the well-ordering of the ordinals)

(K4) There is a smallest ordinal not definable in L—call it σ_K.

But then

(K5) 'is the smallest ordinal not definable in L' defines σ_K.

But then σ_K is definable in L after all! Contradiction.

Berry's paradox: From

(B3) There are natural numbers that aren't 1000-definable in L,

infer

(B4) There is a smallest natural number not 1000-definable in L—call it n_B.

But then

(B5) 'is the smallest natural number not definable in L' defines n_B; and (even after abbreviations are unpacked) this definition has length less than 1000.

But then n_B is 1000-definable in L after all! Contradiction.

According to classical logic solutions to the paradoxes, the problem is in passing from

(K4) σ_K is the smallest ordinal not definable in L

to

(K5) 'is the smallest ordinal not definable in L' defines σ_K.

I explained 'defines' in terms of 'true of', so (K5) amounts to

(K5*) 'is the smallest ordinal not definable in L' is true of σ_K and of nothing else.

This is the conjunction of

(K5*a) 'is the smallest ordinal not definable in L' is true of σ_K

and

(K5*b) 'is the smallest ordinal not definable in L' is not true of anything other than σ_K.

The paradox can be turned into a classical derivation of

$$\neg(\text{K5*a}) \vee \neg(\text{K5*b}).$$

So either (K4) $\wedge$ $\neg$(K5*a), or else (K4) $\wedge$ $\neg$(K5*b).

Either

(K4) $\wedge$ ¬(K5*a):

> σ_K is the smallest ordinal not definable in L, but ‘is the smallest ordinal not definable in L’ is not true of σ_K,

or

(K4) $\wedge$ ¬(K5*b):

> there is a σ other than σ_K such that ‘is the smallest ordinal not definable in L’ is true of σ, even though σ is not the smallest ordinal not definable in L.

(a) involves Underspill

> [F(o), but ¬True(<F(o)>)]

and (b) involves Overspill

> [True(<F(o)>), but ¬F(o)].

The König paradox shows once again that it’s inevitable to have at least one of Underspill and Overspill, given classical logic. (Similarly for the Berry.)

Is this an accurate portrayal of classical solutions?

Some classical theorists say that the notion of definability-in-L shouldn't be taken to be part of L.

I explained 'definable in L' in terms of 'true of', so this requires saying that 'true of' can't be part of L.

> How this could apply to English, which does contain the predicate 'true of', is obscure.

Usually such theorists invoke a hierarchy of restricted truth-of predicates, from which restricted definability predicates are defined.

But *the above analysis then holds for each of the more restricted predicates*.

That is: *for each truth-of predicate in the hierarchy, there is either Underspill or Overspill*.

(Most advocates of a hierarchy of truth-of predicates say that they all involve Underspill only.)

I don't think a hierarchy of truth predicates is an attractive option, but I won't argue that today.

What of the non-classical logic solution to the paradox that I mentioned last time?

On it,

<F(x)> is true of o

is always equivalent to

F(o).

So given how 'defines' was explained in terms of 'true of', there is no room for the difference between

(K4) σ_K is the smallest ordinal not definable in L

and

(K5) 'is the smallest ordinal not definable in L' defines σ_K.

that the classical theorist claims.

Rather, this non-classical solution says that the problem is with the derivation of the existence of a σ_K satisfying (K4).

There is no questioning the derivation of

(K3) There are ordinals that aren't definable in L;

nor of the fact that the ordinals are well-ordered. Even so, it doesn't follow that there is a smallest ordinal not definable in L!

How can this be?

The answer: when we restrict excluded middle, the fact that the ordinals are well-ordered is not properly expressed by the usual least ordinal schema

(C) $\exists\sigma F(\sigma) \vDash \exists\sigma_K[F(\sigma_K) \wedge (\forall\rho<\sigma_K)\neg F(\rho)]$,

but rather by the more general version

(G) $\exists\sigma[F(\sigma) \wedge (\forall\rho<\sigma)(F(\rho)\vee\neg F(\rho))] \vDash \exists\sigma_K[F(\sigma_K) \wedge (\forall\rho<\sigma_K)\neg F(\rho)]$.

Within classical mathematics, (G) reduces to (C), since excluded middle can be assumed to hold generally there.

> (Classical math doesn't contain 'true' or similar predicates, just restricted ones for which excluded middle holds.)

But if we restrict excluded middle for predicates like 'true of', the more general form (G) is all we have when the formula F contains such predicates.

'Definable' was explained in terms of 'true of', so to pass from there being ordinals undefinable in L to there being a smallest one, we have to make a controversial application of excluded middle.

> In particular: we have to assume that for a segment of ordinals whose upper bound is not definable in L, every ordinal in that segment is either definable in L or not.

The paradox shows that this application of excluded middle must be rejected.

That excluded middle must be restricted in application to 'true' is a moral I also suggested last time, in application to the heterologicality and Liar paradoxes.

But this new application is of special interest, for two reasons:

1. The diagnosis of König/Berry is centrally relevant to the "Super-Liar paradox" I mentioned at the end of last time.

> This involves a sentence S that says of itself that for some legitimate iteration D^α of the determinately operator, S is not D^α-true.
>
> Attempt to derive a contradiction from such sentences involves the same fallacy as in the König and Berry paradoxes. (Too big an issue to go into today.)

2. The diagnosis of König/Berry suggests that excluded middle failures *might* be more widespread: **vagueness**.

A standard puzzle about vagueness:

(A) Bertrand Russell was old when he died, at approximately 3×10^{18} nanoseconds of age.

So by the least number principle

(B) There is a smallest natural number N such that he was old at N nanoseconds at age.

Presumably once old he stayed old; and presumably also he wasn't old when he was born, i.e. N>0.

So there's a sharp cutoff point: up through N−1 nanoseconds he wasn't old, but by a nanosecond later he was and remained so.

It is natural to feel that this is contrary to the evident vagueness of the predicate 'old', and that it is highly counterintuitive.

Of course there are ways of trying to argue that (B) shouldn't be counterintuitive, and that it isn't contrary to the vagueness of 'old'. I won't get into that.

What I will do is simply

(i) *raise the possibility* that the logic I've recommended for the semantic and property-theoretic paradoxes might apply more broadly, to vague language;

(ii) consider whether this would have consequences for the issues about rational revisability of logic.

On the application of the logic to vagueness:

In this logic, we can only get from

> (A) Bertrand Russell was old when he died, at approximately 3×10^{18} nanoseconds of age

to

> (B) There is a smallest natural number N such that he was old at N nanoseconds at age,

if we have *an additional excluded-middle premise*: essentially, that at every moment he was either old or not old.

Of course, if you accept classical logic generally, or even for all predicates not involved in semantic or property-theoretic paradoxes, you will accept this added premise. But if you don't, you won't.

But there's more to be said here. In particular, **there seem to be close relations between this approach to vagueness and the approach I've recommended to the semantic paradoxes.**

One way to see the close connection is to return to Lukasiewicz continuum-valued logic.

This has been the logic of choice for most people who have recommended a non-classical logic for vagueness. But it is inadequate, **and for reasons very closely related to the reasons why it is inadequate for dealing with the semantic and property-theoretic paradoxes.**

The main reason is connected with the fact that the Lukasiewicz logic allows for a definition of a determinately operator.

This would be a problem if the determinately operator D obeyed excluded middle, that is, if

$$\mathrm{D}A \vee \neg \mathrm{D}A$$

were a logical law. For in that case, even though 'old' has no sharp boundaries, 'determinately old' does. This would destroy the whole point of the approach.

Fortunately, this doesn't happen: the graph of the determinately operator is as follows:

GRAPH

and so there's a range of values for *A* on the right hand side where it is "fuzzy" whether *A* is determinately true.

So far, then, no problem.

As discussed last time, the determinately operator can be iterated.

Any finite iteration yields a similar graph, but staying at 0 for longer and going up more steeply to 1 near the end, e.g.

GRAPH

But, for any such finite iteration D^n, there is still a range at the end where $\mathrm{D}^n A$ still has values intermediate between 0 and 1, so $\mathrm{D}^n A$ can still be "fuzzy".

So far, still so good.

But what about infinite iteration?

We can define $D^{\omega}A$ as

> For all n, the sentence obtained by placing n occurrences of D before A is true.

In that case, $D^{\omega}A$ has value 1 if A has value 1, and has value 0 otherwise.

So applying D^{ω} to any sentence—or any predicate—turns it into a precise sentence or predicate.

So even though we reject the idea that there's a magic moment at which Russell became old, we must suppose that there's a magic moment at which he became determinately$^{\omega}$-old!

This would make the whole move to restricting excluded middle pointless.

Recall that this is exactly the same point that defeated the continuum-valued semantics as a solution to the semantic paradoxes:

> because of the possibility of defining D^{ω}, we could construct sentences that assert their own lack of D^{ω}-truth, and an analog of the Liar paradox returns.

We saw last time that there are ways to modify the semantics that avoids the problem. (I didn't give the details, just the general idea.)

One gets a semantics that is in many ways like the continuum-valued semantics.

It contains a determinately operator with many of the features of the operator in the continuum-valued semantics.

But that operator doesn't collapse to the 1-0 operator on iteration. Not only doesn't it collapse to it *at* ω, it doesn't collapse to it *however far you go*.

[If you insist on extending it past where our systems of ordinal notations break down, it does in a sense collapse.

But not to anything that could be called an iteration of D: the resultant operator has value less than 1 *even as applied to sentences with value 1*.

Also, it's indeterminate where the breakdown in systems of notations occurs.]

It turns out that exactly the same semantics that works for the semantic paradoxes can be applied to vagueness too.

Indeed, the semantics avoids another feature of the continuum-valued semantics that seems unnatural for vagueness: it avoids the linear ordering of values.

Because of that linear ordering in continuum-valued semantics, sentences like

(*) It is either the case that if Bob is rich then Tim is thin, or that if Tim is thin then Bob is rich

come out as logical truths. But whatever motivation there is for rejecting

(**) Either Bob is rich or Bob isn't rich

would seem to be equally motivation for rejecting (*).

Once again, continuum-valued semantics doesn't quite live up to the motivations.

But the semantics alluded to last time doesn't take the values to be linearly ordered. If it is applied to vagueness, it will not yield (*).

> [There's an alternative presentation of the semantics that is probably more natural for vagueness: in terms of a similarity relation on 3-valued worlds. The conditional is then a "variably strict conditional" of a slightly non-standard sort.]

****

Does this raise a problem for the line I've been taking in these lectures, on the rational revisability of logic?

What I argued last time is that the best logic, in light of the semantic paradoxes, is non-classical.

This is an argument for a rational revision of logic, *if one assumes that the logic we've used up until now is classical*.

But another possible view is that the non-classical logic that my semantics validates is, in some sense, the logic we've been using all along!

And **the fact that the logic seems to work well for vagueness, while somewhat strengthening the case for the logic, also strengthens the case that it's the logic we've used all along.**

If so, I may seem to have **undermined** my case for the rational revisability of logic.

Actually there's a possible worry here even independent of vagueness.

Graham Priest:

> He's suggested that the best solution to the semantic paradoxes is to use a different kind of non-classical logic than the one I've recommended. (A "dialetheic" logic that licenses the acceptance of certain contradictions.)
>
> In some places, he seems to suggest that his line requires supposing that *such a dialetheic logic is the one that ordinary people really employ in practice, before being corrupted by Frege's incorrect codification of that practice*.
>
> **So the suggestion is: *we don't really revise the logic we employ*, we only revise our *theories about* what logic we employ.**

That would solve the puzzle with which I began these lectures.

> It isn't puzzling how we can rationally revise *our theories about the logic we employ*, they can be revised in just the ways that our theories about everything else can.
>
> It is rational revision in *the logic that we employ* that can seem problematic, because it would be a revision in a very fundamental norm, and the logic must itself be used in any rational argument for revision.

But Priest's picture (as given in these passages) seems dubious.

When Frege codified classical logic (and when Priest, Lukasiewicz and others have codified various non-classical alternatives), they haven't primarily been theorizing about what logic people employ, but formulating candidates for a logic *to* employ.

> [Frege attack on psychologism: dealing with the "laws of thought", but only in a normative sense.]

What logic people in fact employ isn't directly relevant. (If it's a clear question at all.)

Admittedly, a proposal to employ a logic that's at total variance with how people actually reason would have little point. But that doesn't undermine the distinction between a recommendation to use a certain logic as our all purpose logic and a theory about ordinary practice.

Still, it's not out of the question that a proposed alternative to classical logic *does* correspond to how ordinary people reason.

And in the case of vagueness, it does seem that a logic without excluded middle does *capture something in ordinary thought*.

> I think the ordinary speaker does have some inclination to resist the conclusion that there is a first nanosecond at which Russell was old, and to resist the conclusion that at some given moment he either was or wasn't old.
>
> This isn't to say that they accept a non-classical logic rather than a classical one. (I'm skeptical of there being a clear answer to the question of what their logic is.) Only that **they don't *clearly* accept classical logic.**

But then, what becomes of my case for a rational revision of the logic *that we employ*? Why shouldn't we say that people who are puzzled by the semantic paradoxes are simply making what are mistakes *by the very logic they employ*?

> E.g., they don't realize that 'true' is a vague term, and they should reason with it in the more general way that they do with other vague terms;
>
> Or, they do realize this, but make what are mistakes *by their own light* in their reasoning.

Before getting to the main response, an important subsidiary.

I've implicitly made a 3-way distinction:

> first, the logic we employ;
>
> second, our theory about what logic we employ;
>
> and third, our view about what logic we ought to employ.

A question about all of these—especially the first two—is what is meant by "employing a logic".

I partly addressed this in the second lecture. There I argued that to employ a given logic is, roughly, to accept its dictates for degrees of belief:

> if $A_1, ..., A_n$ obviously entail B in a given logic, then errors of reasoning aside, one who employs that logic will believe B to at least degree $\Sigma P(A_i) - (n-1)$.

But I now want to focus on the hedge "errors of reasoning aside". (The errors here are both of omission and of commission.)

Attributing a logic to a person is a matter of idealizing their practice—deciding which pieces of reasoning involved errors of commission or omission.

But what constitutes a good idealization? What distinguishes accepting a logic but erroneously reasoning in it on a given occasion, from accepting a different logic on which that "erroneous reasoning" was legitimate?

One picture of idealization involves the competence-performance distinction:

> there is a logic that in some deep sense governs our epistemic behavior, but
>
> various performance errors (inattention, memory-limitations, and the like) interfere with it.

On such a model, this "deep logic" would be "the logic we employ".

I'm skeptical of that model. But skepticism about it breeds skepticism about the clarity of "the logic we employ". Attributing a logic to a person is a matter of idealizing their practice, *and idealizations needn't be unique*.

The possible lack of uniqueness is far less relevant to the third question distinguished above ("What logic *should* we employ?") than to the others ("What logic *do* we employ?" and "What is *our theory about* what logic we employ?"). The third is in many ways the clearer question.

Back to the worry about vagueness.

The worry was:

> 1. Ordinary people don't, at the most basic level, reason in accordance with excluded middle when it comes to vague terms.
>
> Maybe they've been corrupted by Frege and others into thinking that they do or that they ought to. And maybe this sometimes leads them to actually do so. But if so, that's a kind of performance error.
>
> 2. The semantic paradoxes all turn on excluded middle, or on other principles that are suspect in absence of excluded middle.
>
> 3. Moreover, the reason these principles are inapplicable to semantic terms like 'true' is that those terms have a kind of vagueness (or at least, something closely akin to it) once they are applied outside of their "safe" range.
>
> 4. Ordinary people either don't realize the vagueness in semantic terms or else misapply their own logic of vagueness, and that's why they are taken in by the paradoxes.

A substantial part of this can be freed of its uncritical use of "the logic ordinary people employ":

> As long as people have *some* tendency to reason without excluded middle with vague terms, we get: there isn't a *clear* shift of logic in accepting a non-classical solution to the paradoxes.

There's something to this. But for two reasons, it doesn't undermine my case for the rational revisability of logic.

1. It's somewhat plausible that the ordinary person resists some applications of excluded middle and the least number principle;

BUT: it doesn't seem at all plausible that the ordinary person employs a logic of the conditional that avoids semantic paradox.

It's a very delicate matter to get a logic of the conditional that avoids semantic paradox while capturing a decent part of ordinary reasoning. It seems extraordinarily unlikely that any remotely realistic idealization of ordinary practice would lead to such a logic.

It's doubtful that any remotely realistic idealization of ordinary practice would lead *even to a logic adequate to vagueness*.

> (From a conditional one can define a determinateness operator which can then be iterated into the transfinite. Avoiding collapse to an operator that would trivialize vagueness requires a lot of delicacy.)

And avoiding semantic paradox involves still tighter constraints on the conditional.

So: **Even if one postulates that "the ordinary person's logic" is non-classical, it almost certainly will require revision to accommodate the paradoxes.**

And the process of revising it to do so seems eminently rational.

2. Speculations about "what logic the ordinary person employs" seem beside the point.

Even if we suppose that the ordinary person does in some deep sense "employ" a given non-classical logic that adequately handles the paradoxes, we can easily imagine people that don't.

Couldn't they, on learning of logics that do handle the paradoxes compatibly with the intersubstitutivity of True(<*A*>) with *A*, be persuaded that it is best to alter their inferential practice to bring it into accord with such a logic?

That they could is, in effect, what I tried to argue last time.

Sociological speculation that the ordinary person already at some deep level employs such a logic seem not only

dubious

but also

irrelevant to the philosophical point.

**

Should be time left:

Could catch-up on stuff from Lecture 2 on truth-preservation issue, to show that what's changed is norms of reasoning, not views on what preserves truth.

Or Curry paradox.

Or something on modal semantics.

Prepare slides on last two. (To insert at an appropriate place.)

Modal semantics

This involves a set W of 3-valued "worlds". A parameterized formula gets one of the values 1, ½ and 0 at each world.

The connectives ¬, ∧, ∨, ∀ and ∃ are treated in the way you'd expect (strong Kleene semantics) within each world.

For the conditional, we need a similarity structure on each world: we assign to each world w a system F_w of (non-empty) "spheres of similarity around w".

> In dealing with counterfactual conditionals, Lewis assumed that the spheres around any world were nested, which says in effect that for any worlds u and v, either u is more similar to w than v is, less similar, or equally similar.
>
> I employ a weaker assumption ("directedness"): that for each w and any two spheres U and V around it, there is a sphere contained in the intersection of U and V.

(Without the stronger assumption, 'neighborhood' is a better term than 'sphere'.)

The clause for ‘→’:

The value of A→B at w is

- **1 if there is a neighborhood around w throughout which the value of A is less than or equal to that of B;**
- **0 if there is a neighborhood around w throughout which the value of A is greater than that of B;**
- **½ otherwise.**

[In ordinary vagueness cases, we can assume that each world is in all the neighborhoods around it. In applying the semantics to the paradoxes, we need to allow for “abnormal worlds” where this condition isn’t met. But the boldfaced method for evaluating the conditional is the same in both cases.]