What is analyticity?

(*First written in 2019*)

In "Sense and Reference", Frege defines a statement as *analytic* if and only if it is true in virtue of the meanings of the terms. (Given his context, this fact has immense historical interest, which unfortunately I can't discuss here.) But what is it for a statement to be "true in virtue of the meanings of the terms"? In this post, I sketch a theory that I favour. I don't claim that it's original (it probably isn't). It's also very un-Fregean.

First off, what are the desiderata? Historically, the candidates have been: (i) that analyticity is a property of sentences, (ii) that it explains their truth regardless of what states of the world obtain, and (iii) that it explains why we have a priori knowledge of the statements expressed by analytic sentences.

But is it wise to seek a concept that satisfies (i) - (iii)? I think that history has shown that it isn't. In retrospect, it now seems unlikely that we're going to be able to circumscribe some special, infallible domain of knowledge that's explained by our understanding of our own language, since understanding a language is oftentimes fallible and a posteriori. So I'm not going to approach this topic as an over-ambitious epistemologist who is seeking for a key to unlock the foundations of knowledge (like the philosophers of the past). Rather, I approach this topic as a philosopher of language who thinks that there's an intuitive difference between the sentences "all triangles have three angles" and "all ravens are black." Pretheoretically, there's a difference in how we should explain their truth values, and it has something to do with how language works.

Background Assumptions

I need to expose my theoretical presuppositions. I'll simply list them.

(i) Truth is primarily a property of *propositions,* and only derivatively a property of *sentence tokens.*

(ii) A sentence token is true if and only if it expresses a true proposition.

(iii) I give a two-tiered view of meaning that is reminiscent of Peter Strawson's distinction between 'meaning' and 'proposition' and David Kaplan's distinction between 'character' and 'content.' I will call the two levels 'conventional meaning' and 'content.' Here's how I understand them.

(iii.a) *Conventional meaning* is a property of *sentence types* and *expression types* individuated by *patterns of use in a linguistic community.* Informally, it can be thought of as a product of the *conventions of language use* that govern the use of the expression types (and for sentences, it is determined compositionally). Theoretically, it plays the role of being an important determiner of the content of a sentence *token* when used in context. Conventional meaning is the community-wide conventions that determine what an expression token may be used to say (that is, what content it may express) when used in context. It thus includes the Kaplanian character of indexical terms. These conventions are implicitly understood by competent language users. By this, I mean that they have a sort of practical knowledge of how the conventions governing an expression type function to help determine the content of a token when

used in context. That is, they know *how* to use an expression to *say something.* This knowledge is *practical*—it is *not* a type of propositional knowledge. It is more like knowing how to ride a bike or tie one's shoes as opposed to knowing that something is true.[1]

(iii.a) Content is the proposition expressed by a sentence token in its context. It is *what is said* by the sentence token when it is uttered. Theoretically, content has the roles of *bearing truth value, bearing modal properties, being the semantic value of a sentence token in context, and being the objects of the intentional mental states* (what is believed, desired, thought, etc.). It is determined by many factors, including the expression type's conventional meaning, the context of use, the intentions of the speaker, the causal history of the vehicles of content, and perhaps others.

(iv) *Millianism.* I believe that content is individuated as per the view of Mill and Russell. Propositions should be conceived as complexes of objects and properties that are the referents of the sub-sentential parts of the sentence token. Less metaphorically, this means that they should be individuated by (i) reference (in particular, they are more coarse-grained than Fregean propositions which are complexes of senses), and (ii) structure (in particular, they are more fine-grained than Lewisian sets of possible worlds). This presupposition is controversial (albeit, it's the majority view of philosophers in general and philosophers of language in particular); one day I should write something in defence of it, but I won't do that here.[2]

Analyticity

I have two more assumptions: (1) I take it that analyticity is a property of *sentence types;* and (2) *aprioricity* is a modifier of a subject's knowledge of a proposition. (Actually the second one may need to be qualified, for it seems that a subject can come to know one and the same proposition *in a way that's a priori* and *in a way that's a posteriori*. So perhaps it's better to say that apriocity is a modifier of *the way that a subject comes to know a proposition.*)

Now that I've outlined my theoretical background, I can proceed to some positive views on analyticity. But before doing so, let's pause to reflect on why this project is worthwhile. Why can't we just say that *'S' is analytic if it's true in virtue of meaning?* The reason, as should be clear from above, is that there isn't a single notion of meaning to plug into this formula. Rather, there are various types of 'meaning', such as *content* and *conventional meaning.* And it is worthwhile thinking about how these interact to make some sentences true. Unfortunately, there's no straightforward way of doing this. That is because there are at least three different types of analyticity that need to be treated separately. They are: (I) the theorems of propositional logic, and what I will call (II) "Kaplan analyticity" and (III) "Frege analyticity." I thus construct my account in three stages.

I. *The truths of propositional logic.*

---

[1] Update: in retrospect, this notion may be too vague to carry the theoretical weight that I demand of it. If I wanted to make this theory work, this would be the thing to shore up. (Also, I now question the wisdom of labelling this a kind of 'meaning.')

[2] Update: grahamsethmoore.wordpress.com/…/fregeanism-vs-direct-reference-theory

The first stage concerns sentence types whose instances are true in virtue of being a theorem of propositional logic. Sentences like "Either it is snowing or it is not snowing" and "If it is *not not* snowing then it is snowing" are perfect examples of this kind.

I think that there's no getting around the fact that these sentences owe their truth to the *content* of the logical operators "if… then", "not", "and", and "or." These sentences are true because (i) "if… then", "not", "and", and "or" each express their own truth function, and (ii) the syntax of these sentences corresponds to composition of these truth functions into a function that always returns 'true' for any argument.

This explanation of logical truth is fine as far as it goes, but it already creates some prima facie trouble for my background orientation. That is because, firstly, there is good reason to treat the theorems of propositional logic as analytic. (Eventually we're going to want to explain the analyticity of the sentence type "if *a* is a vixen then *a* is a female fox" by its correspondence to the propositional theorem "if *a* is a female fox then *a* is a female fox". So the explanation of the analyticity of the former will depend on the analyticity of the latter.) But this conflicts with my earlier claim that *analyticity is a property of sentence types,* and *content is a feature of sentence/expression tokens.* These two tenets seem to rule out explaining analyticity by appealing to *content*. But I just said that the explanation of the truth of the propositional logical theorems appeals to the content of the logical operators.

I think we're going to have to make an exception for propositional logic. For the sentence types that are theorems of propositional logic, their analyticity will be explained by the content of their logical operators. But luckily, this makes sense. That is because, plausibly, the expression types "if… then", "not", "and", and "or" are *context-invariant* in the English language: each token of them will express the same truth-function in any context. As a result, it makes sense to talk about *the* content (truth function) expressed by the expression-*type* "or", because it will always be the same.[3]

So to summarize, the theorems of propositional logic are analytic, and their truth is explained by the content (i.e. truth functions) of the logical constants. Notice though, that I haven't attempted to explain *why* they are analytic. At this point, I'm regarding it as a stipulation that each of these truths is analytic. In subsequent stages I will build my account of analyticity by treating the theorems of propositional logic as base cases.

II. *Kaplan analyticity*

Having made this one concession to appeal to the content of the propositional logic constants, from now on I intend to appeal to only conventional meanings to explain analyticity. This invites the following account:

Analyticity. A sentence type 'S' is analytic if and only if the conventional meanings of its singular terms, predicates, quantifiers and operators, combined with the content of its propositional logical vocabulary, suffice, on their own, to explain why any proposition expressed by its tokens (in a single context) will be true.

The type of analyticity that this analysis is best suited to explain is known in the literature as *Kaplan*

[3] Afterthought: this is an idealization (really a fudge); there are plenty of uses of "or" in English that don't express disjunction; e.g. "You can have tea or coffee."

analyticity. Kaplan analytic sentences are those whose *characters* (i.e. functions from contexts to contents that are part of conventional meaning) are such that they return a true content for any context. (Characters, recall, are one type of conventional meaning that pertain specifically to indexicals.) The paradigmatic example of such a sentence is "I am here now." Although this sentence type will express different propositions when used in different contexts, the proposition that it will express in any context will always be a true one. (Other plausible examples of Kaplan-analytic sentences include the T- and R-sentences (e.g. ""Snow is white" is true iff snow is white", ""Socrates" refers to Socrates",[4] and the formal truths about knowledge ("if S knows that P then P").

There are some interesting features of Kaplan analyticity that overturn conventional thinking about analyticity. In particular, the *propositions expressed* by a Kaplan analytic sentence are typically *contingent* and *a posteriori* (recall "I am here now"). Moreover, Kaplan analytic sentences may fail to express a proposition when used in context. So for example, ""Socrates" refers to Socrates" will fail to express a proposition if the token of 'Socrates' on the right hand side fails to refer (say, if I intend to refer to someone who doesn't exist). Nonetheless, the sentence *type* is analytic. (Note that this possibility doesn't cause any problem for my definition of analyticity.) The upshot is that, even when a sentence type is analytic, one may need additional empirical investigation to know (i) whether its token expresses a proposition at all, and (ii) if it does, which proposition it expresses.

II.ii *first-order logic and beyond*

One might wonder why I only discussed propositional logic in (I), but didn't say anything about quantificational logic or modal logic. The answer to this is because the quantifiers and modal operators are *context sensitive,* and so their conventional meanings do not return constant contents in every context. If I utter, in one context, "every beer is in the fridge", and I utter in a different context "every beer is in the fridge", I may be talking about different sets of beer. On that token, if I utter "every beer is a beer" in the first context, and "every beer is a beer" in the second context, then although I'm guaranteed to have said something true in both contexts, the true propositions I express could be distinct. Which proposition I express will depend on which 'domain of objects' is supplied by context. Similar remarks reply to modal sentences.

The upshot is that we should treat the sentence types that are quantificational and modal theorems along the lines of Kaplan-analytic sentences. As such, the fact that they always express true propositions is traceable to the *characters* of the quantifiers and modal operators. The characters of the quantifiers and modal operators may be thought of as functions from contexts (which supply domains of quantification or the sets of relevant possible worlds) to contents. It is part of their conventional meaning of these sentences that, given any domain of quantification or set of possible worlds supplied by context, their tokens will always express true contents.

There is a complication that affects our readiness to *know* that a token of these sentence types says something true. Namely, the context may subtly change mid-utterance, allowing for the quantifiers embedded in the sentence to answer to different domains of quantification. This possibility means that an *analytic sentence type* may express a falsehood. Note though, that this doesn't rescind its status as

[4] The explanation for this is rather complicated. See Ori Simchen, "Meaningfulness and Contingent Analyticity, *Nous* (2003).

analytic: an analytic sentence type will express a true content (if it expresses a content at all) when tokened *in a single context.* But if the context shifts, then all bets are off. Here's an example. Suppose that Lewisian contextualism is true about knowledge ascriptions, and then consider the following sentence type: "My evidence rules out all of the possibilities where I do not have hands; but if my evidence rules out all of the possibilities where I do not have hands, then my evidence rules out all of the possibilities where I am a handless brain-in-a-vat; therefore, my evidence rules out all of the possibilities where I am a handless brain-in-a-vat." First-order logic would tell us that this is a valid argument, and that its conditionalization will always express a truth *when the domain of quantification is constant*. But the contextualist contends that the domain is *not* constant, and so its conditionalization may express a falsehood.[5] The contextualist also notes that average speakers might not be privy to these subtle context shifts. This is yet another reason why there is no straightforward path between *analyticity* and privileged knowledge.

III. *Frege Analyticity*

My account so far has held up well, albeit it has had some surprising conclusions. But the most challenging obstacle for a Millian account of analyticity has to do with sentences that might be called "Frege analytic." These are examples such as "all bachelors are unmarried" and "all vixens are female foxes."

The reason why these are challenging for the Millian is precisely because the Millian claims that content is identical to reference. So "all vixens are female foxes" expresses the proposition <for all x in domain D, if x is V then x is V>, which is a trivially true proposition. *However,* the sentence "all water is H20" *also* expresses a similarly trivially true proposition for the Millian, and yet it *shouldn't* count as analytic. The upshot is that our explanation of the analyticity of the former should not also entail the analyticity of the latter, and so we cannot appeal to content alone.

Frege thought that the solution was to abandon Millianism about content. According to him, content should be fine-grained. So whereas "vixen" has the same content as "female fox", "water" does not have the same content as "H20." We can then explain the analyticity of the former by saying that, when the term "vixen" is substituted for its synonym "female fox" we yield a logical truth, but we can't likewise yield a logical truth from "all water is H20" (two terms are synonymous when they share the same content).

But if we're not going to follow Frege on this (and I don't believe we should), we're going to have to find an alternative explanation as to why "all vixens are female foxes" is analytic whereas "all water is H20" isn't. Clearly we're going to have to appeal to elements of meaning other than content. Our explanation cannot rely on the claim that "vixen" and "female fox" are 'synonymous.'[6]

This is admittedly the sketchiest part of my proposal, but I do think that this can be solved. Here is my proposal in a nutshell: I believe that it is part of the English conventions surrounding the expression type "vixen" that it has a descriptivist *metasemantics.* To see what I mean, recall that in *Naming and*

[5] In this case, one is a contextualist about 'evidence.' See Jonathan Ichikawa, *Contextualising Knowledge,* OUP, 2017.

[6] Some Millians (e.g. Nathan Salmon) will bite the bullet and concede that "all water is H20" is analytic afterall, despite our inability to know this truth from mere linguistic competence. I think it's preferable to avoid this conclusion.

Necessity, Kripke distinguishes between two types of descriptivism. On one type (the semantic version), it is claimed that a name is *synonymous* with a description (i.e. has the same content as a description). On the other type (the metasemantic version), a description plays the role of *fixing the referent* of the name, even though the name isn't *synonymous* with the description. Kripke argues against both types of descriptivism for *most* natural language names. But that leaves open the possibility that a select few expressions have a descriptivist metasemantics.

Let's conjecture that the content of a token of the type "vixen" is fixed as whatever satisfies the description "female fox." (This explains *how* the content is fixed, but it doesn't identify *what* the content *is*.) If true, then this fact would be part of the conventions of language use for the expression type "vixen", and its role would be to determine the content of the tokens of "vixens." In short, the descriptivist metasemantics is part of the *conventional meaning* of the expression type "vixen."

If I'm right about this, then it's easy to see that the sentence type "all vixens are female foxes" would satisfy my analysis of analyticity given above. That is because the metasemantics of the term "vixen" will be such that a token of it is guaranteed to express the same content as a token of the term "female fox", *regardless of what the content is.* In other words, as a matter of the *conventional meaning* of "all vixens are female foxes", it is guaranteed to express a proposition that has the form <all Fs are F>, in a way that is predetermined *prior* to what F turns out to be.

There would be no such guarantee for the sentence "all water is H20." That is because, plausibly, the metasemantics of the term "water" is given by something like the ostension story told in Putnam's "Meaning of Meaning", and is independent of the description "the substance composed of H20." So to determine that "all water is H20" expresses a truth, one has to actually attend to the particular proposition expressed.

Since my account of analyticity doesn't appeal to the alleged synonymy of "vixen" and "female fox", it is fairly unorthodox. However, my account has two virtues that a synonymy-based account lacks. First, my account nicely explains an observation made by Kit Fine. The observation is that, when we say that "all vixens are female foxes" is 'true by virtue of meaning' we intuitively say that it's true by virtue of *the meaning of the term "vixen.*" In the explanation of the sentence's analyticity, there is a sort of asymmetry between the terms "vixen" and "female fox"; the term "vixen" is the most central to the explanation. My account accommodates this by placing the burden of the explanation on the *metasemantics* of the term "vixen." The synonymy account has no explanation for this asymmetry.

The second advantage of my account is that it seems to make better sense of Twin Earth cases. Imagine a Twin Earth scenario where the speakers of Twin-English have all of the same intersubjective conventions regarding the expression type "vixen"—*including the convention to fix its referent via the description "female fox"*—but instead of foxes, their planet is inhabited by another species that's ostensibly similar. If the Fregean is going to be a good externalist, then even they must admit that the English utterance "all vixens are female foxes" and the Twin English utterance "all vixens are female foxes" express different contents. For this reason, the Fregean explanation of their analyticity must be different for each case.

But intuitively that's unnecessary. Pretheoretically, if "vixen" has the same intersubjective conventions of use on both Earth and Twin Earth*,* then both analyticities should have the same explanation *despite their difference in content.* My account makes good on this; the Fregean account

doesn't.[7]

Conclusion

Let me wrap up by restating the main features of my account and its funny, unorthodox consequences. The most basic idea of my account is that analyticity is not ultimately attributable to *what you say* (except for when you use the propositional logical terms "or", "and", "if… then" and "not"). Rather, it is explained by the conventions surrounding the form of words you choose to *express whatever it is that you say.* The idea is that: sometimes these conventions conspire to predetermine that *whatever you say, it's going to be true* (provided that you've managed to say something at all, and you haven't changed the context mid-sentence). More precisely, my account was given by:

Analyticity. A sentence type 'S' is analytic if and only if the conventional meanings of its singular terms, predicates, quantifiers and operators, combined with the content of its propositional logical vocabulary, suffice, on their own, to explain why any proposition expressed by its tokens (in a single context) will be true.

And now for the surprising consequences. Throughout this essay, I've remarked that:

- An analytic sentence may express a contingent proposition; e.g. "I am here now."
- An analytic sentence may express a proposition that isn't knowable a priori; e.g. "I am here now."
- An analytic sentence may fail to express a proposition at all; e.g. "'Socrates' refers to Socrates" in the context of reference failure.
- *An analytic sentence may even express a falsehood*; e.g. in the case where there are multiple tokens of the same type of context-sensitive expression, and the context changes mid-sentence. E.g. the conditionalization of the following argument, according to contextualists: "My evidence rules out all of the possibilities where I do not have hands; but if my evidence rules out all of the possibilities where I do not have hands, then my evidence rules out all of the possibilities where I am a handless brain-in-a-vat; therefore, my evidence rules out all of the possibilities where I a handless brain-in-a-vat." For a simpler example, consider "Now is the same time as now", but spoken *verrrryyy* slowly.

Afterthought from the year 2022. I would be remiss not to mention one of the most unorthodox consequences of them all. This account of analyticity is tailor-made to pertain *only* to the expressions of *public language.* With the exception of the truths of propositional logic, it has nothing to say about the analyticity of *thoughts* (e.g. beliefs). You might think that the *belief* that *all vixens are female foxes* is analytic, but the *belief* that *water is H20* is not. Well, this account has nothing to explain this difference. Indeed, on this account, it would be a category error to ascribe analyticity to beliefs, since analyticity is reserved for *expression types*. I'm currently inclined to accept this consequence and concede that there

[7] This argument assumes that the best version of Fregeanism involves externally-individuated, *de re* senses. A 'narrow content' version of Fregeanism could evade this criticism; but I just don't think that such a view is plausible.

isn't such a thing as *conceptual truth* for mental contents. Here we see how radical the departure is from the traditional notion of analyticity—and how useless this notion is for epistemological purposes.