**Scott Soames**

# Cognitive Propositions, Truth Functions, and the *Tractatus*

It has been 141 years since, first Frege, but then Russell and Wittgenstein, put the metaphysics and epistemology of propositions at the center of philosophy of language, logic, mathematics, and mind. Whereas Frege gave us the first great system for understanding language and thought, his metaphysics and epistemology of propositions was unrealistically other-worldly.[1] The same was true of the early Russell who, by 1912, had become dissatisfied with his inability to explain the intentionality of the abstract structures he had heretofore called "propositions" – which were at best artificial models of what we really assert, believe, and know.[2] It was then that he was struck with the idea that what unites the elements of assertion, belief, and hypothesis, and gives them representational content are the minds of agents, without which neither truth nor falsity can be understood. Although this led him to the dead end that was his multiple relation theory of judgement,[3] Wittgenstein was more successful in putting a human face on representational thought. The tractarian picture theory of propositions we use sentences to express was grounded in the use of linguistic and other artifacts to represent items in the world as being one way or another.[4] Despite remaining submerged for nearly a century, this seminal idea has now been put in cognitive form and developed by several leading philosophers.[5]

The central tenets of the new cognitive approach include T1-T3.

T1 Propositions are neither abstract *know-not-whats,* nor merely instrumentally useful models to be interpreted away in extracting empirical content from semantic theories; rather they are real cognitive acts, operations, or products of such that are central to the mental lives of human beings and many animals.

---

[1] Frege ([1879] 1967, [1892], 1970)
[2] Russell (1912) chapter 12. See also Soames (2010) chapter 4 and Soames (2014) chapter 9 section 2.
[3] See Soames (2014) chapter 9 sections 3 and 4.
[4] See Soames (2016, 2018 chapters 2 and 3).
[5] See King (2007), King, Soames, and Speaks (2014), Soames (2015), Hanks (2015), Jesperson (2010, 2012, 2015), and Moltmann (2017).

T2 The cognitive relations we bear to propositions when we judge, believe or know to be *x,y, or z* to be *such and such or so and so* are as explicable and commonplace as our cognitive relations to x, y, or z, or to their sizes, shapes, colors, or other properties,

T3 Our acquaintance with propositions, and our ability to predicate properties of them, are fundamental aspects of human self-consciousness.

In addition to making a naturalistic epistemology of propositions possible for the first time, the metaphysical identification of propositions with spcific types of cognitive acts or operations has greatly expanded the solution space for the many variations of what has arguably been the most recalcitrant nest of empirical problems in semantic theory since its inception, Frege's puzzle.

This is the context in which I here return to the *Tractatus Logico Philosophicus,* which was one of the founding texts that inspired my own view of cognitive propositions. I start with Wittgenstein's obscure remark at 3.12 that an elementary *proposition is a propositional sign* [a sentence] *in its projective relation to the world.*[6] My reconstruction of it follows the *Tractatus* in rejecting the Frege-Russell conception of propositions as abstract objects the intentionality is explanatorily prior to, and independent of, that of the cognitive agents who passively entertain them. Like Wittgenstein*,* I take elementary propositions to be, or to be abstracted from, meaningful uses of sentences. Such uses are repeatable cognitive act types, performances of which are events in which agents use sentences in accord with the semantic conventions that govern them. My task here is to explore the feasibility of extending this account to truth-functionally compound propositions in a way that captures something approximating the putative tractarian virtue of *identifying truth-functionally equivalent propositions,* while avoiding the tractarian vice of identifying necessarily equivalent propositions.

Wittgenstein motivates this putative virtue in part by saying that 'not', 'and', and 'or' don't denote anything. To get a feel for this, think of simple cases involving predicates formed by prefixing the copula to conjoined, disjoined, or negated general terms -- e.g., *is red and red, is*

[6] Suggested in Soames (2016, 2018 chapter 2).

*red or red,* and *is not not red.* It is appealing to think that the property each is used to predicate is simply *being red.* Next consider conjunctions and disjunctions of the sentence *a is red* with itself, and double negations of that sentence. Using 'and', 'or', and 'not' to form compound sentences rather than to form compound general terms or predicates may seem to simply repackage the earlier result. Of course, even this doesn't give us the conclusion that propositions expressed by truth-functionally equivalent sentences are identical. There is, however, a not implausible argument that seems to. Although in the end I will show that, it doesn't succeed, the steps along the way illuminate and extend the conceptual resources of the cognitive conception of propositions.

## I: The Model

Consider a tiny plane-geometrical universe containing squares, circles and triangles each of which is entirely red, green, or yellow. I will use three simple first-order languages to describe it. Each is partially, but not entirely, tractarian. The discussion that follows makes much of an important feature of the conception of propositions as purely representational cognitive act types. It follows from this conception that propositions can be cognitively distinct by virtue of imposing different conditions on minds that entertain them, while being representationally identical in the sense that they represent the very same things as being the very same way. The tractarian considerations investigated below raise the question of whether a new species of cognitively distinct but representationally identical propositions should be recognized, and also whether some cognitive demands placed on understanding a sentence are not parts of the individuation of cognitive propositions.

The non-logical vocabulary of my languages consists entirely of Millian proper names, (at least one for each object) and general terms ‘red’, ‘green’, ‘yellow’ designating properties *being red, being green,* and *being yellow* (true of all and only red, green, and yellow things). Similar

remarks apply to 'square', 'circular' and 'triangular'. Since I will need a distinction between essential and non-essential properties, I will here stipulate that, for our purposes, each of the geometrical properties will count as essential to each particular that has it, whereas the color properties won't. A further, more general stipulation is added in order to keep the calculations needed to keep track of things as simple as possible. The three color predicates will be assumed to be mutually exclusive and disjunctively exhaustive. Thus in all possible states of my tiny universe, each object will bear exactly one of those colors. (We don't exclude colors transparently definable from these colors -- e.g. *being either green or yellow* -- though there will be no simple terms designating them.). This could be modified; we could add a colorless predicate, but in what follows I will simply exclude colorlessness. A corresponding stipulation will govern the shape properties.

**II: Do truth-functionally compound predicates express the same property?**

My first language, L1, also contains functors 'and', 'or', and 'not' that combine with general terms to form general terms. They designate functions that map properties, or pairs of such, expressed by their arguments, to properties expressed by the correspondingly complex general terms. Thus there is no upper bound on the number of general terms – 'red', ' (not red), '(red or green)', '((red or green) & circular)', etc. The same is true of predicates, each of which consist of the copula plus a general term. Atomic sentences consist of a name plus a predicate.

'Not' maps the property designated by its argument onto its complement, 'and' maps pairs of properties onto the corresponding intersective properties while 'or' maps them onto "additive" properties. *Being not red* is true of all and only the circles, squares, or triangles that are either green or yellow. *Being green or yellow* is itself a color, called 'grellow' in the metalanguage; it is true (at any world-state w) of all green, and also of all yellow, circles, squares, triangles (at w) and nothing else. This property, *being grellow,* is metaphysically simple, just as *being red, being*

*green,* and *being yellow* are. It is not composed of *being green, being yellow, or being red* plus any metaphysical building block corresponding to 'or' or 'not'. It is metaphysically simple, even though different ways of expressing it -- e.g., 'not red' and 'green or yellow' – are complex. Thus, *not being red = being green or yellow*. Similarly *not being square = being circular or triangular.*

The same is true of *being red, being red or red, being red and red, and not being not being red* (or *being not not red* for short). Since *not being red = being green or yellow,* and *being not not red = being red,* it follows that *not being green or yellow = being red.* Ditto for the geometrical properties. Next consider *not being red* and *not being square.* Since the former = *being yellow or green* (i.e. *grellow*) and the latter = *being circular or triangular* (i.e. *circtangular*), we see that *not being red and not being square = being grellow and circtangular,* while *not being red or not being square = being grellow or being circtangualar.* (We are, of course, speaking in the metalanguage here.) Thus, *not being (grellow or circtangular) = being red and square,* and *not being (not red and not square) = being (grellow or circtangular).* For any general terms P and Q, *being (not P) and* (*not Q) = not being (P or Q)* and *being (P or Q) = not being (not P and not Q).* In short, *truth-functionally equivalent predicates* in L1 express the same properties.

Call an atomic sentence of L1 simple or complex depending on whether its predicate is. Each complex atomic sentence is equivalent to infinitely many truth-functional compounds of simple sentences (expressed in an extension of L1 containing sentential negation, conjunction, and disjunction). For example, the complex atomic sentence *A is (square and not red),* is associated with truth-functionally complex sentences, including *A is square and* (*A is yellow* or *A is green*). Two complex predicates of L1 are *truth-functionally equivalent iff* their associated truth-functional compounds of simple atomic sentences with the same subjects are truth-

functionally equivalent. This is the understanding we invoke when we say that *truth-functionally equivalent predicates* of L1 express the same properties. Next we identify something that I will provisionally call *the atomic proposition expressed by a use of sentence S of L1 in accord with the linguistic conventions governing S.* It is a cognitive act type of predicating the property expressed by the predicate of the referent of the name. It is true iff the object designated by the name has the property expressed by the predicate.

We now have, for each name designating an object and each predicate expressing a property, infinitely many sentences which, when used in accordance with the linguistic conventions governing them, predicate the same property of the same thing, and hence express the same atomic proposition, despite placing different cognitive demands on their users. Note, I do *not* say that this is the *only* proposition expressed by a use of any of the given sentences. I say that there is one such proposition that is commonly expressed by uses of all them.

**Table 1**

*A is red, A is not,not red, A is red and red, A is red or red,…*

*A is red or square, A is not (not red and not square), A is not ((green or yellow) or (circular or triangular)), …*

*A is (red and not square), A is red and (circular or triangular)), A is (not (green or yellow) and not square), A is not (square or (green or yellow)), A is [(red and circular) or (not (green or yellow) and not (square or circular))]…*

The sentences in each set predicate the same property of the same thing. Thus, there is a simple atomic proposition p such that each of the sentences expresses p. The calculations needed in the first four sentences to bring the property to mind and predicate it of the referent of 'A' are trivial. But that isn't true across the all the sentences of all the sets. This prompts two conclusions.

First, the cognitive acts of *using the various sentences* in a single set to predicate the same property of the same thing get more complex as one moves through the sequence of sentences. If

we abstract away from this complexity by not taking it to enter into the identity conditions of the propositions expressed by uses of these sentences, we will then take each sentence in the same set to express a single proposition, and indeed the same proposition in each case. For the first set, it will be the unique proposition that *a is red,* understood as the cognitive act of predicating redness of a – no matter the routes by which *redness* and *a* are brought to mind. Similarly for the other sets. Assuming, as is standard, that there are only finitely many names in the language, this means that only finitely many propositions are expressible in the language, even though there are infinitely many sentences.

However, we could conceptualize things differently. Consider the cognitive steps used in bringing to mind the properties expressed by compound predicates as we move through the sentences on a given line of table 1. If we include those steps in the identities of propositions predicating those properties, then the different sentences on a single line will express propositions that are *representationally identical but cognitively distinct* from the cognitive act type that abstracts from how the relevant property predicated of an object is brought to mind. In set 1 of table 1 this proposition is the act of predicating *being red* of *a* -- no matter how *a* and *redness* are brought to mind (including by using other sentences, or without using sentences at all). On this alternative conception, to use one of the sentences following the first is to entertain a cognitively more complex proposition, the entertainment of which automatically counts as also entertaining the cognitively simplest proposition that *a is red,* which we have at the first step. Similarly for the other sets. On this second conception, the cardinality of expressible propositions matches that of the sentences that express them.

Although this inflationary conceptualization is coherent, it also seems extravagant, at least from the point of view of the *Tractatus.* There, propositions are taken to be abstract "logical" pictures, of which sentences are perceptual signs. The interest in propositions –- which it is the

function of sentences to express -- is simply in what they represent.[7] Thus, for Wittgenstein, propositions that represent the same things as being the same ways are identical.[8] The idea is powerful, even apart from the *Tractatus.* Propositions are cognitions the essential nature and function of which is simply to represent the world, or parts of the world, as being certain ways. As long as what we wish to represent doesn't include the cognitive acts of ourselves and others, it is natural to take propositions to be identical iff they are representationally identical. That will change when we turn our attention to the propositional attitudes of agents, but for the moment the metaphysically deflationary conception of propositions seems preferable to the inflationary one.[9]

Whatever we decide about these two ways of individuating cognitive propositions, we may want to resist the idea that all sentences in the same set have the same meaning. We will resist, if we tie *knowing the meaning of S* closely to *understanding S,* since (i) to understand S one must understand its vocabulary and its grammatically significant constructions, and (ii) we are often reluctant to accept the claim that two sentences differing only in the substitution of computationally or conceptually equivalent words, phrases, or sub formulas (e.g. in arithmetic) must mean the same thing, if, as is often the case, it is possible for someone to understand both without being able to recognize that they are true of the same things (e.g. natural numbers). Think of this as a (potential) difference between the semantic content of a sentence and its meaning -- one that has nothing essential to do with indexicality or non-indexicality.

Before we go further, let us take stock. Our language L1 is resembles the envisioned ideal language of the *Tractatus* in four respects: (i) it's names are Millian, (ii) its simplest sentences are used to predicate properties (or relations) of the objects designated by the names in the sentence, (iii) truth-functional operators on properties -- 'and, 'or' and 'not' -- don't denote any

[7] The relationship between sentences and propositions in the *Tractatus* is discussed on pp. 36-39 of Soames (2018).
[8] See sections 3.1 and 3.2 of Soames (2018).
[9] The tractarian view of propositional attitudes is critically discussed on pp. 91-96 of Soames (2018).

(non-cognitive and non-linguistic) items in the worldly reality that sentences are used to talk about, (iv) truth-functionally equivalent sentences differing only in their "logical" vocabulary express the same proposition. That said, the two languages do, of course, differ in many ways. Among the most important differences, for our purposes, are: (a) that tractarian names, but not names in L1, refer to so-called metaphysical simples, (b) that the individual atomic sentences of the *Tractatus,* and sets of them, are logically and metaphysically independent; this is not true of simple atomic sentences L1, and sets of them, and (iv) that necessarily equivalent sentences of the *Tractatus* are truth functionally equivalent, and express the same proposition, whereas that is not so for the sentences of L1. For example, ignoring questions of existence, if both *A is square* and *B is circular* are true, then they are necessarily equivalent, without being truth-functionally equivalent, and without expressing propositions that have anything to do with each other.[10]

Next we will repackage these results in language L2, which has the same names and the same simple predicates as L1, but a more conventional syntax. As before, the language will contain 'and', 'or', and 'not', but this time they will operate on sentences to form compound sentences, rather than on predicates to form compound predicates. Nevertheless, the sentential operators will share a basic similarity with the predicative operators in L1.

**III: Do truth-functionally equivalent compound sentences express the same proposition?**

Consider conjunction and disjunction in L2. Given a pair of sentences 'A is red' and 'B is square' – expressing the propositions *that a is red* and *that b is square,* (i) we map those propositions onto the properties *being such that a is red* and *being such that b is square,* each of which is either true of everything or true of nothing, (ii) we conjoin/disjoin the properties, giving us *being such that a is red and/or b is square,* and (iii) extract the relation R, *being such that _ is red and/or _ is square,* which we predicate of *a,b* (in that order). This is the proposition *that a is*

[10] Though the examples are more complex we can get essentially the same result without ignoring existence (despite having no existence predicate and no quantifiers in L1).

*red and/or b is square* expressed by 'A is red and/or B is square'. It represents *a* as standing in R to *b* (and nothing further). The negation of this proposition – the proposition *that it's not the case that a is red and/or B is square* – is similar; it predicates ~R of *a,b* (and nothing further). It represents *a* as not standing in R to *b*.

In moving from 1-place to 0-place properties (and relations extractable from them), we have not changed the metaphysics of properties (and relations). For any n-place (including 0-place) properties P, Q, their negations/conjunctions/disjunctions are metaphysically simple properties that are true of n-tuples that the original property isn't true of/that the original properties are both true of/that one or the other of the original properties are true of. Thus, our construction seems to suggest that the conjunction or disjunction of any proposition with itself should be identical with the original proposition. The conjunction/disjunction of the proposition *that a is red* with itself predicates a certain property of a -- the property *being such that a is red and/or a is red* – which, it is natural to think, is just the property *being such that a is red.* What is the relationship between predicating *being red* of *a* and predicating *being such that a is red* of *a*? Although they are different cognitive acts, it is hard to resist the idea that they are representationally identical. Both merely represent *a* as being red.

Supposing this to be so, consider the higher-order act of *representing a as red*, which can be performed in either of these two ways (among others). Perhaps this act is the ur-proposition *that a is red* – the purely representational cognitive act that is performed no matter whether (i) an agent simply brings *a* and the property *being red* to mind, predicating the latter of the former, or (ii) an agent also generates the 0-place property *being such that a is red* and predicates it of *a*, or (iii) an agent goes further, generating the pair *being such that a is red, being such that a is red,* conjoins/disjoins the two, and predicates the result of *a* (all of which are different lower-level acts), and so on. Note, this higher-order act – this ur-proposition – is not identified with any

particular act of predicating a property of an object. It is a more abstract than that. It is act that is performed by performing any one of many specific acts of predication (each of which, in these cases involves predicating *redness* of *a*) – just as the act of traveling to work is an act that is always performed by performing a more specific act of traveling – e.g. of driving, walking, biking, or to work.

What about the double negation of a proposition? Suppose one starts by bringing the object *a* and the property *being red* to mind, predicating the latter of the former. From this one may form the 0-place property *being such that a is red.* Negating this property gives us *not being such that a is red*; predicating it of *a* gives one the proposition *that a is not such that a is red*. Repeating the process, one predicates *not being such that a is not such that a is red* of *a*. This property is, it seems, identical with the 0-place property *being such that a is red.* Predicating it of *a* gives one the proposition *that a is such that a is red,* which is an instance of the abstract ur-proposition *that a is red,* just as before.

What we have done is, in effect, to produce the L2 version of the first line of table 1 above. Versions of other lines could similarly be reproduced. What about sentences like (1-3)?

1. ~(a is red or (b is circular or b is triangular))

It is, of course, equivalent to

2. *~ (a is red) and ~(b is circular or b is triangular),*

which, adhering to the assumptions of our simple model, is equivalent to

3. (a is green or a is yellow) and (b is square).

If (1-3) express the same propositions as (1λ-3λ), the results for L1 should extend to L2.

1λ. λxy ~(x is red or (y is circular or y is triangular)) a,b
2λ. λxy (~ x is red and ~(y is circular or y is triangular) a,b
3λ. λxy (x is green or x is yellow) and (y is square) a,b

The move from 2λ to 3λ depends only our simple metaphysics of properties plus the L1 account of operations on properties. As for the move from 1λ to 2λ, the representational content of the

latter is the same as the representational content of the pair of L1 sentences *a is not red* and *b is not (circular or triangular).* Thus, it would seem, (1-3) might be taken to express the abstract ur-proposition that *a is green or a is yellow and b is square* instances of which involve quite different sequences of predications. If so, then we should be in a position to conclude that truth functionally equivalent sentences of L2 that are neither tautologies nor contradictions will express the same abstract proposition. Although this may seem attractive, we are not, in fact, yet entitled to this result. To see this, we must first investigate propositional attitudes, after which we will consider an extension, L3, of L2 that contains attitude ascriptions. Only then will we be able to come to a principled decision about whether truth-functionally equivalent sentences of L2 that have the same non-logical vocabular express the same propositions.

## IV: Cognitive Propositions and Propositional Attitudes

We begin with ascriptions of the form (4).

4. X asserted/believed, knew that S

A sentence of this form is used by a reporter R to represent an agent as standing in the relation of asserting/believing/knowing to one or more propositions expressed by R's use of S. I say "one or more" because uses of sentences often express multiple cognitive propositions.[11] Suppose, in the simplest case, that the embedded sentence S is one in which a name, indexical, or variable relative to an assignment is used to designate an individual *a*, and a predicate is used to predicate a simple property P of *a*. In such a case, the semantic content of S relative to a context is, typically, the proposition – *Bare Prop* – that merely represents *a* as being P, without imposing any constraint on how either the object or property are brought to mind.

Since propositions are complex cognitive acts that involve *sub acts* of identifying objects and properties -- which agents always identify *in some way* -- agents always entertain bare propositions by entertaining cognitively enriched but representationally identical versions of

[11] Soames (2015), chapters 3-5.

them. The enrichments – modes of Millian, rather than Fregean, presentation -- are particular ways of bringing objects and properties to mind (and predicating the latter of the former). These ways include, but are not limited to, identifying objects and properties perceptually (via visual, auditory, or other sense experience), identifying them linguistically (using specific words, phrases, etc.) and identifying them via *de se* and/or present-tense cognitions. Abstracting away from these, we reach cognitively less demanding propositions the entertainment of which places *no restrictions on how the properties or objects are brought to mind*. By entertaining the enriched proposition we, thereby, entertain the bare proposition too.

One of the most significant Millian modes of presentation is first person cognition, the logic of which is illustrated by (5) and (6).

5a. I am in danger. *Said by SS*
b. SS is in danger.

6a. I believe that I am in danger. *Said by SS*
b. SS believes that SS is in danger.

Since (5a) and (5b) express representationally identical but cognitively distinct propositions, (6a) can be false even if (6b) is true. This happens when I see SS in a mirror and believe him to be in danger, without believing I am in danger. Here, we distinguish predicating property P of an agent SS cognized *in the 1st-person way* from predicating P of SS however cognized. To do the first is to do the second, but not conversely, so the acts are different. Since the same property is predicated of the same agent, the propositions are cognitively distinct but representationally identical . In this way, we capture the fact that my epiphany--*I am the one in danger*--involves believing a truth I hadn't previously believed, even if my believing it is just my coming to believe, in a new way, something already believed.[12]

[12] See Chapter 3 of (Soames 2015) plus the discussions of Perry, Lewis, Kripke, and Kaplan in chapter 9.

Propositions about the attitudes agents bear to propositions are often sensitive to cognitive differences among representationally identical propositions stemming from the way agents are taken to cognitively identify predication targets in the proposition that serves as the object of the attitude. To see this, let 'logicism' be a Millian name for the proposition L, *that arithmetic is reducible to logic*, which may be designated by a *that*-clause in an attitude ascription. Since L is what the two terms contribute to the representational contents of (7), sentences (7a) and (7b) express representationally identical, but cognitively different truths.

7a. Russell tried to prove (the proposition) *that arithmetic is reducible to logic*.
b. Russell tried to prove *logicism*.

Entertaining, accepting, asserting or believing these proposition requires predicating *trying to prove* of Russell and L. In the case of (7a) one is also required to identify L as the target of predication by entertaining it. This cognitively enriched proposition is the semantic content of sentence (7a). By contrast, the semantic content of sentence (7b) does not constrain how one identifies L when predicating *trying to prove* of Russell and L. Here, any form of acquaintance with L will do – including either naming or entertaining it. In addition to this bare proposition, there is also a cognitively enriched, but representationally identical proposition, that *requires* one who entertains it to identify L using the name 'logicism'. This proposition is one of the propositions that a speaker who assertively utters (7b) would normally assert, along with the corresponding bare semantic content of (7b). Normally, however, this wouldn't add much, since very little, in any, information over and above that present in the bare proposition would be gained. This is, of course, not true when one considers moving from the bare proposition to the cognitive enrichment required by (7a), which tells one a great deal more about what Russell was trying to prove.

Because of this difference between (7a,b), (8a,b) can differ sharply in cognitive significance, as well as truth value.

8a. Mary believes that Russell tried to prove that arithmetic is reducible to logic.
b. Mary believes that Russell tried to prove logicism.

If Mary picked up the name 'logicism' merely by hearing it used to designate some thesis in the philosophy of mathematics that Russell tried to prove, (8b) may be true, even if she hasn't heard what he thought about arithmetic, in which case (8a) will be false. This can be so only if the two attitude ascriptions represent Mary as believing different things, and so are representationally distinct. The truth of (8a) requires Mary to believe a proposition that can be entertained and accepted only by one who entertains L, and thereby acquires information about what Russell was trying to do. Because she doesn't entertain L, (8a) is false. Since she does believe a related proposition that represents Russell as trying to prove L, (8b) is true, even though she remains in the dark about what he was up to.[13]

With this in mind, consider the true, representationally identical, but cognitively distinct propositions expressed by sentences (9a-c).

9a. Logicism is (the proposition) that arithmetic is reducible to logic.
b. Logicism is logicism.
c. That arithmetic is reducible to logic is (the proposition) that arithmetic is reducible to logic.

Anyone who entertains, accepts, or bears a relevant attitude to the semantic content of (9a), must identify the second argument, L, of the identity relation by entertaining it; in the case of (9c) that requirement applies to both arguments, while in (9b) it applies to neither. Now notice two further facts: First, anyone who entertains, believes, or knows proposition (9c) counts as entertaining, believing, or knowing propositions (9a) and (9b). Second, since proposition (9c) is knowable apriori, it follows that the other two propositions are too – in the sense that there is *a way of having them in mind* that allows one to see their truth, without requiring empirical justification.

[13] Ibid., pp. 39-43 and 80-8.

Thinking back about Mary, that may seem surprising. After all, if she knows *that Russell tried to prove logicism*, and if it is knowable apriori *that logicism is the proposition that arithmetic is reducible to logic,* then shouldn't Mary be able to conclude, by reflecting on what she knows, *that Russell tried to prove that arithmetic is reducible to logic*? Of course not. The reason she can't is similar to the reason why Kripke's puzzling Pierre can't conclude that New York is pretty, despite the fact that is an apriori consequence of a pair propositions both of which, we may imagine, he believes.[14] The propositions are (i) that London is pretty, which he sincerely assents to when it is presented to him by the sentence 'Londres est jolie' of this native French, and (ii) the proposition that New York is pretty if London is, which he believes to be true by falsity of antecedent when presented to him by the English sentence.[15]

The case is similar to instances of Frege's Puzzle involving 'Hesperus' and 'Phosphorus'. Unlike most names, these names really do conventionally encode different specific information. Nevertheless, they are Millian in the standard sense. Like other names, the semantic contents of these two names are their bearers, which is what they contribute to the compositionally determined semantic contents of all sentential clauses in which the names occur, including those governed by modal operators (e.g. *necessarily, possibly).* In both cases, what they contribute is simply their bearer, Venus. What is unusual about 'Hesperus' and 'Phosphorus' is that there is more to *understanding* them than simply being able to use them to designate their bearer. One who uses 'Hesperus' and 'Phosphorus' is expected know that users typically presuppose that 'Hesperus' stands for something visible in the evening and 'Phosphorus' stands for something visible in the morning. One who mixes this up *misunderstands* the names. Those who know enough to use the names know this.

With this in mind consider (10).

[14] Kripke, "A Puzzle About Belief."
[15] I here pretend that indicative conditionals in English are material conditionals. We could, of course, reconstruct the case with a real material conditional.

10a. Hesperus is a planet.
b. Phosphorus is a planet.
c. x is a planet. (with Venus as value of 'x')

Let proposition p be expressed by (10c). $P_H$ is a proposition representationally identical to p that requires one to identify the predication target, Venus, of *being a planet* via the name 'Hesperus'. $P_P$ requires cognition via 'Phosphorus'. Utterances of (10a) assert both $P_H$ and p; utterances of (10b) assert $P_P$ and p. Hence, one who accepts (10a) may believe something different from what one believes in accepting (10b), thereby explaining the potentially different truth conditions of utterances of (11a,b), even though the propositions believed are representationally identical.

11a. Mary believes that Hesperus is a planet.
b. Mary believes that Phosphorus is a planet.

Now consider A's use of (12), addressing B, each presupposing that both *understand* the names.

12. Hesperus is Phosphorus

A asserts not only the bare singular proposition that predicates identity of Venus and Venus, but also the cognitively distinct, but representationally identical, proposition entertainable only by identifying Venus via the two names. Although this proposition merely represents Venus as being Venus, B extracts more information from A's assertive utterance. Presupposing that A *understands* the names, B reasons that A knows A will be taken to be committed to the claim that *the unique object that is both Hesperus and visible in the evening is the unique object that is both Phosphorus and visible in the morning*. Knowing that A expects him to so reason, B correctly concludes that A *asserted* the descriptively enriched proposition (as well as the unenriched, bare proposition).

Note, this proposition is contingent, even though A's assertive utterance of (13) is true.

13. In fact, it is necessarily true that Hesperus is Phosphorus.

How can this be? The extra representational content carried by A's assertive utterance of (12) arises from the linguistically enhanced proposition asserted, the presupposition that A and B understand the names, and the information that comes with this understanding. Although these same factors are present when A utters (13), they do *not* lead to comparable enrichment. The difference between the enrichment of A's use of (12) and the lack of such enrichment of A's use of (13) hinges on what *understanding* 'Hesperus' and 'Phosphorus' requires. *It requires knowing that most agents who use them take, and expect others to take, 'Hesperus' to stand for something (actually) seen in the evening and 'Phosphorus' to stand for something (actually) seen in the morning*. Presupposing that both and B understand the names in this sense, A and B add descriptive content to A's utterance of (12). Since taking the names to refer to things *actually* seen at certain times tells one *nothing* about when they are seen at other *possible* world-states, A and B *don't* descriptively enrich occurrences of the names under the modal operator when evaluating assertive utterances of (13).

This explanation depends on three points. First, to cognize o via a name n does *not* involve *predicating being named n of o* -- any more than cognizing oneself in the 1st-person way involves predicating that one is so-cognized. Second the linguistically enhanced propositions asserted by utterances of sentences containing names are representationally identical to, but cognitively distinct from, the bare semantic contents of the sentences uttered. Third, to *understand* an expression requires not only the ability to use it with its semantic content, but also the knowledge and recognitional ability needed to use it to communicate with others in ways widely presupposed in the linguistic community. This dynamic extends to natural kind terms, where it provides solutions to many instances of Frege's puzzle involving them.[16]

[16] See chapter 4 of Soames (2015)

## V: Transparent vs. Non-Transparent Representational Identity

In section 2, I presented a simple language L1 and a simple semantic model to go with it. Its only linguistic complexity involved compound predicates gotten by negating, conjoining, and disjoining predicates. This was combined with a metaphysical picture of properties that encouraged the thought that those expressed by "truth-functionally equivalent" *compound predicates* (containing no Millian singular terms) are identical. Section 3 traded negations, conjunctions, and disjunctions of predicates for negations, conjunctions, and disjunctions of sentences in a language L2, using the mechanism of 0-place properties. An argument was then offered for identifying *propositions* expressed by truth-conditionally equivalent *sentences* of L2.

The strategy employed used lambda abstraction to turn compound sentences of L2 – e.g., *'~ a is green and ~ a is yellow'* – into simple sentences – *'λx (~ x is green and ~ x is yellow) a'* – which, by the assumptions of the model, could be given maximal reductions – e.g. *'a is red'*. The argument was extended to identify propositions (1-3). using the supposedly proposition preserving steps (2aλ) and (2bλ) to get from (2) to (3)

1. ~(a is red or (b is circular or b is triangular))
2. ~ (a is red) and ~(b is circular or b is triangular)
3. (a is green or a is yellow) and (b is square)

2aλ. ~ (a is red) ➔ λx [~ (x is red)] a
2bλ. ~(b is circular or b is triangular) ➔ λy [~(y is circular or y is triangular)] b

We are now in a position to identify a serious flaw in this reasoning. Suppose we take it for granted that propositions (1) and (2) can be identified – i.e. that both are instances of the abstract proposition that represents *a as not being red* while representing *b as not being circular* and also representing *b as not being triangular*. We assume, in other words, that performing the predications involved in understanding sentences are different ways of entertaining the same abstract proposition, identified by its representational content (without imposing further

constraints on the cognitive route to that content). What we can't show, given this, is that proposition (3) is identical with proposition (2) (or proposition 1).

To see this, it is helpful to consider the transition from (2) through (2aλ) and (2bλ) to (3). First, (2aλ) together with the assumptions of our model (illustrated in L1), tells us that for object a not to be such that a is red, is for a not to be red, which is for a to be *green or yellow*. Next we appeal to the assumptions of our model, which tell us that the property *not being circular or triangular* is the property *being square*. Thus, a proposition that predicates *λy [~ (y is circular or y is triangular)]* of b represents b as being square (and nothing more). So, if we accept (2bλ), we will think that ~ (*b is circular or b is triangular)* does too, which may tempt us to take propositions 2 and 3 to be identical. Although this seems natural, it is refuted by a well-known objection.[17]

Imagine a group of ancients who, when they looked at Venus in the evening sky named it 'Hesperus', and who, when they looked at Venus in the morning sky named it 'Phosphorus' -- not realizing that they had named the same thing twice. Next imagine that today a fully competent speaker of English, John, says something about Hesperus and Phosphorus – i.e. Venus -- by assertively uttering sentence (14a), which he understands.

14a. The ancients believed (and asserted) that 'Hesperus' referred to Hesperus and 'Phosphorus' referred to Phosphorus.

In so doing, John asserts, and expresses his belief in, the true proposition that is the semantic content of sentence (14a), which is also the semantic content of sentence (14b).

14b. The ancients believed (and asserted) that 'Hesperus' referred to Venus and 'Phosphorus' referred to Venus.

In moving from (14a) to (14b), we are not saying that John didn't also assert and believe other propositions too, including one requiring Venus to be cognized once via the name 'Hesperus'

[17] The objection is presented in Soames ([1987] 2009). A prominent objection is raised in Edelberg (1994), and refuted in Soames ([2008] 2009).

and once via the name 'Phosphorus'. In the right sort of context, this proposition may have been one of the propositions John asserted and believed. But since doing that also counts as asserting and believing proposition (14b), there is no blocking the move to (14b).[18] It is tempting at this point to invoke lambda abstraction, giving us (14c).

14c. The ancients believed (and asserted) that $\lambda x$ ('Hesperus' referred to x and 'Phosphorus' referred to x) Venus.

But since (14d), which is false, is trivially derivable from (14c), the latter must also be false, which means that in invoking lambda abstraction we have moved from a truth to a falsehood.

14d. The ancients believed (and asserted) that 'Hesperus' and 'Phosphorus' were coreferential.

The problem isn't, of course, with lambda abstraction per se, but with applications of it that collapse two or more occurrences of a Millian name in different argument places of an n-place predicate into a single occurrence providing a single object as argument of a related n-1 place predicate. Frist consider (15a,b,c).

15a. 'Hesperus' referred to Venus and 'Phosphorus' referred to Venus
b. λx ('Hesperus' referred to Venus and 'Phosphorus' referred to x) Venus.
c. λxy ('Hesperus' referred to x and 'Phosphorus' referred to y) Venus, Venus

If the complement clause (15a) of (14b) is replaced by (15b,c), no problem results; truth is preserved and (14d) is not derived. It is only the move to (15d) that is problematic.

15d. λx ('Hesperus' referred to x and 'Phosphorus' referred to x) Venus.

Why is it problematic? If we ask ourselves, *What things are represented as being what ways?* It is arguable that, despite the different structures of these sentences, the propositions they express all represent each of the two names 'Hesperus' and 'Phosphorus' as referring to Venus (while representing nothing further]. Thus, it is hard to deny that that propositions (15a-d) are *representationally identical,* and so have identical truth conditions. Nevertheless, the propositions differ in how informative they are to agents who come to accept them. Whereas

[18] Soames (2015) pp. 18-19, 22-25, 73-79, 157-58.

entertaining. and accepting proposition (15d) requires recognizing Venus as the *common referent* of both 'Hesperus' and 'Phosphorus', entertaining and accepting propositions (15a-c) sometimes provides an agent with no way of recognizing this.[19] This is so when the agent cognizes Venus in different ways – e.g. via different Millian names, via different visual perceptions, or once via a name and once via a perception – as arguments of predications.

This difference between propositions (15a-c) and (15d) is, of course, cognitive. Although the former are (arguably) representationally identical to the latter, they are not *transparently* so. Any way of cognizing and accepting proposition (15d) provides an agent who accepts it with a way of cognizing and accepting propositions (15a-c), but the converse is not true. We may put this by distinguishing two relations – transparent vs. non-transparent -- representational identity. Propositions P and Q are transparently representationally identical *only if agents* who entertain and accept one of them are always, thereby, in a position to derive and accept the other – as they are with (15a-c). By adding a similar condition on their negations, ~P and ~Q, we can define an equivalence relation; *Propositions P and Q* are transparently representationally identical *iff* agents who entertain and accept one are always, thereby, in a position to accept the other, and similarly for their negations.

## VI: Identity and Truth-Functional Equivalence Revisited

The argument in section II failed to establish that truth-functionally equivalent sentences of L2 express the same proposition because it failed to distinguish transparent from non-transparent representational identity. It may well be that all those sentences – including (1-3) are representationally identical (given the assumptions of our model), but once the cognitive nature of propositions is recognized, it becomes clear that this is not enough. A proposition imposes not only conditions the world must satisfy, if it is to conform to how it is represented to be, but also

[19] These are cases in which the agent doesn't recognize the recurrence of the entity Venus in the agent's act of entertaining the proposition (or, sometimes, of recurrence of the same name in a sentence used to express it. See chapters 6-8 of Soames (2015) for discussion.

conditions a mind must satisfy, if it is to entertain the proposition. What we have just rehearsed is an often neglected subcase of a much larger reality that cognitive propositions were designed to address – representationally identical propositions that are cognitively distinct.

The subcase involves pairs of representationally identical propositions p and q, in which the entertainment of p, but not q, requires distinct cognitions of an object playing different argumentative roles, without further restrictions on the modes of presentation of those objects. (Think of these as *bare propositions* in the sense of section III.) When these propositions are expressed by sentences, sometimes the objects are the referents of different, but coreferential, Millian names, and sometimes they are the referents of different occurrences of the same name. The same point can be made using occurrences of demonstratives, variables relative to assignments (needed for cases of quantifying-in), or natural kind terms. Despite the limited nature of our linguistic examples, this cognitive reality is ubiquitous.[20]

It is important not to confuse the cognitive lesson to be learned with the claim that the structure of propositions must parallel the structures of the sentences expressing them. Nothing we have seen requires us to distinguish the bare propositions expressed by sentences (15a-c) despite their structural differences. Their semantic contents could be the same – the act of representing Venus as the referent of each of the two names (and nothing further), without placing any constraints on how Venus is cognized. On this view, understanding and accepting sentences (15a-c) would all count as entraining that single abstract proposition. So would understanding and accepting (15d). However, the proposition that is the semantic content of (15d) would add the requirement that Venus be recognized once as the referent of the two names. Such a distinction is sensitive to certain sentence structures, even though it is not defined in terms of structure.

[20] See section 4 of Soames ([1987] 2009).

Once this is understood, we can delineate the extent to which the conclusions drawn regarding the language L1 might be expanded to L2 and to L2λ, which adds lambda abstraction to L2. First Let S1 and S2 be any pair of logically contingent truth-functionally equivalent sentences of L2 that have no multiple occurrences of Millian terms (i.e. no multiple occurrences the same Millian term, or of different terms with the same Millian content). Nothing I have said prevents us from taking the propositions semantically expressed by these sentences, as well as any sentences generated by lambda abstractions to be transparently representationally identical to one another. Thus propositions of this sort, which don't further restrict how relevant objects and properties are cognized, may well be identical. However, further than this we cannot go.[21]

Finally, I return to a suggestion broached in section 2 that the conjunction or disjunction of a proposition P with itself is identical with P. The initial, seemingly compelling, idea was that the original proposition was representationally identical with the conjunction or disjunction, *without involving any significant cognitive difference*. To vindicate this idea we have to show that the representational identity of the bare propositions expressed by sentences (16a,b,c) is transparent.

16a. a is a planet
b. a is a planet or a is a planet
c. a is a planet and a is a planet

In fact, the representational identity is non-transparent. Suppose Sam refers to Venus when he uses 'Hesperus' or 'Phosphorus'. We may imagine he is confident that either 'Hesperus' or 'Phosphorus' refers to a planet, but he is unsure which. Then he may entertain and accept the bare proposition (16b) expressed by *'Hesperus is a planet or Phosphorus is a planet'*, without thereby being in a position to accept proposition (16a). Hence the representational identity of propositions (16a) and (16b) is not transparent. The representational identity of (16b) and (16c) is similarly non-transparent. The non-transparent representational identity of propositions (16a)

[21] Unless we allow the Millian mode *recognition of recurrence* to distinguish cognitively distinct but representationally identical propositions, as discussed in Soames (2015), chapters 7,8.

and (16c) follows from the non-transparent representational identity of their negations (16~a) and (16~c).

16~a. ~(a is a planet)
16~c. ~ (a is a planet and a is a planet)

The only remaining question concerns negation. Do cognitive propositions provide a framework in which the double negation of a proposition P can be identified with P? Given the assumptions of the simple model, I see no reason why not. At any rate, no doubt has been cast on the argument given in section II. We begin by bringing the object *a* and the property *being red* to mind, predicating the latter of the former. From this we form the 0-place property *being such that a is red.* Negating it gives us *not being such that a is red*; predicating it of *a* gives us the proposition *that a is not such that a is red*. Repeating the process, we predicate *not being such that a is not such that a is red* of object *a*. Now compare (17a,b) with (18a.b), where 'n' and 'm' are coreferential Millian names for object a

17a. n is not such that n is not such that a is red.
b. λx [x is not such that x is not such that a is red] n.
18a. m is not such that m is not such that *a is red*.
b. λx [x is not such that x is not such that *a is red*] m.

Given the nature of these structures, the agent knows that their representational contents are the same, no matter what 'n' and 'm' refer to. Each merely represents a as being red (and nothing more). Since this is precisely what is represented by simply predicating redness of a, there is nothing preventing us from taking all these acts to be instances of the more abstract act that is the proposition *that a is red.* This is what is meant by saying that proposition p and proposition ~ ~p are identical. I don't say we have proven that p and its double negation are identical. I do say that the idea that they are identical is a significant tractarian holdover that is consistent with the simple cognitive model developed here.

## References

Edelberg, Walter (1994)"Propositions, Circumstances, and Objects," *Journal of Philosophical Logic 23.*

Frege, Gottlob ([1879], 1967) *Begriffsscrift* trans. by S. Bauer-Mengelberg, in Van Heijenoort, Jean, ed. 1967. *From Frege to Gödel.* Harvard University Press.

_____, (1892], 1970) "On Sense and Reference," trans. by Max Black, in Geach and Black, eds. *Translations from the Philosophical Writings of Gottlob Frege* 1970.

Hanks, Peter(2015) *Propositional Content*, New York: Oxford University Press.

King, Jeffrey C., (2007) *The Nature and Structure of Content*, Oxford: Oxford University Press.

King, Jeffrey C. and Scott Soames and Jeff Speaks (2014) *New Thinking about Propositions,* Oxford: Oxford University Press.

Jesperson, Bjorn (2010), "How Hyper are Hyper Propositions?" *Language and Linguistics Compass,* 39: 296-30.

_____, (2012)"Recent Work on Structured Meaning and Propositional Unity," *Philosophy Compass*, 7:620-30.

_____, (2015) "Proliferation of Propositions," *Thought*

Moltmann, Friederika (2017 ) "Attitudinal Objects and the Distinction Between Actions and Properties," in Friederika Moltmann and Mark Textor, eds. *Act Based Conceptions of Propositional Contents*, Oxford: Oxford University Press

Russell, Bertrand ([1912] 1997 ) *Problems of Philosophy,* Oxford University Press.

Soames, Scott([1987], 2009) "Direct Reference, Propositional Attitudes, and Semantic Content," Philosophical Topics, 15; reprinted in Soames *Philosophical Essays Vol. 2.*

_____ ([2008] 2009) "Why Propositions Can't be Sets of Truth-Supporting Circumstances," *Journal of Philosophical Logic* 37, reprinted in Soames *Philosophical Essays Vol. 2*

_____ (2010) *What is Meaning?,* Princeton University Press.

_____ (2014), *The Analytic Tradition in Philosophy*. Vol. I, Princeton University Press

_____ (2015), Rethinking Language, Mind, and Meaning, Princeton University Press.

_____ (2016) "Propositions, the Tractatus, and "The Single Great Problem of Philosophy," *Critica*, vol. 48, No. 143, 2016, 3-19.

_____ (2018), *The Analytic Tradition in Philosophy,* Vol. 2, Princeton University Press.