

THE SEMANTIC FOUNDATIONS OF PHILOSOPHICAL ANALYSIS

SAMUEL Z. ELGIN

“The business of philosophy, as I conceive it, is essentially that of logical analysis”

Bertrand Russell

Abstract. I provide an analysis of sentences of the form ‘To be F is to be G ’ in terms of exact truth-maker semantics—an approach that identifies the meanings of sentences with the states of the world directly responsible for their truth-values. Roughly, I argue that these sentences hold just in case that which makes something F also makes it G . This approach is hyperintensional and possesses desirable logical and modal features. In particular, these sentences are reflexive, transitive, and symmetric, and if they are true, then they are necessarily true, and it is necessary that all and only F s are G s. I motivate my account over Correia and Skiles’ [11] prominent alternative and close by defining an irreflexive and asymmetric notion of analysis in terms of the symmetric and reflexive notion.

§1. Introduction. The subject of this paper is a targeted reading of sentences of the form ‘To be F is to be G ’, which philosophers often use to express analyses and which have occupied a central role in the discipline since its inception. Examples that naturally lend themselves to this reading include:

1. To be morally right is to maximize utility.
2. To be human is to be a rational animal.
3. To be made of water is to be made of H_2O molecules.
4. To be even is to be an integer divisible by 2 without remainder.
5. To be a béchamel is to be a roux with milk.

As these examples indicate, although philosophers frequently utter these sorts of sentences, they do not fall within the exclusive purview of philosophical inquiry. Mathematicians, chemists, and even chefs use them as well. And while some are knowable a priori, others are sensitive to empirical investigation—so there is variation within the phenomenon I seek to describe.

Sentences of this form have been employed since antiquity (as witnessed by 2). Throughout the ensuing history, proposed instances have been advanced and rejected for multitudinous reasons. On one understanding, this investigation thus has a long and rich history—perhaps as long and rich as any in philosophy. Nevertheless, explicit discussion of these sentences in their full generality is relatively recent (for example, see [10–12, 32, 34]). Current advances in hyperintensional logic provide the requisite resources to analyze these sentences perspicuously—to provide the analysis of analysis

Received: September 9, 2020.

2020 *Mathematics Subject Classification*: 03A05, 03B47.

Key words and phrases: analysis, identification, generalized identity, truth-maker semantics.

(or, as some might prefer, the identity of identity). The significance of this project ought to be apparent: the standards that putative analyses must meet hang in the balance.

A bit loosely, I claim that a sentence of this form is true just in case that which makes it the case that something is *F* also makes it the case that it is *G* and vice versa. There is a great deal to say about what I mean by ‘makes it the case that’. In some ways, this paper can be read as an explication of that phrase. For the moment, suffice it to say that rather than understanding it modally (along the lines of ‘To be *F* is to be *G*’ is true just in case the fact that something is *F* necessitates that it is *G* and vice versa), I employ truth-maker semantics: an approach that identifies the meanings of sentences with the states of the world exactly responsible for their truth-values. It will take time before my account can be stated any more precisely; the details of truth-maker semantics must first be appreciated.

The structure of this paper is as follows. In Section 2, I isolate the targeted reading of ‘To be *F* is to be *G*’ at issue. In Section 3, I discuss current developments in truth-maker semantics in considerable detail. The various ways in which the semantics might be constructed bear on philosophical analysis in that the logical features of analysis turn on which alternative we select. For example, the truth-value of sentences of the form ‘To be *F* is to be both *F* \wedge *F*’ depends upon a choice-point in truth-maker semantics—or so I will argue. In Section 4, I provide the details of my account and demonstrate that it has appropriate logical features; it is transitive, reflexive, and symmetric and has the resources to distinguish between the meanings of predicates with (classically) logically equivalent extensions (sentences of the form ‘To be *F* is to be both *F* and *G* or not *G*’ are typically false). I present the resulting logic of analysis and highlight advantages my proposal has over that defended by [11], the closest account available in the literature. I conclude in Section 5 by sketching an account of an irreflexive and asymmetric notion of analysis in terms of the reflexive and symmetric notion. Where possible, I present this view informally; however, there are limits to how informal any systematic discussion of truth-maker semantics can be. Those interested in a more technical overview of the underlying semantics would do well to read [23, 24].

§2. Generalized identity. ‘To be *F* is to be *G*’ may admit of multiple readings. There may be a reading of ‘To be a politician is to be indebted to one’s constituents’, ‘To be a scientist is to be curious about the natural world’, or ‘To be early is to be on time’ on which these sentences are true. Perhaps closer to my intended target, there may also be a reading of ‘To be a bachelor is to be male’ or ‘To be crimson is to be red’ on which these sentences are true. If so, these are not the readings I address. The reading of ‘To be *F* is to be *G*’ that I am concerned with is synonymous with (or, at the very least, close to synonymous with) ‘To be *F* just is to be *G*’ or, perhaps, ‘Being *F* is *the same as* being *G*’.

This reading has borne multiple labels in the literature. Some refer to these sentences as ‘generalized identities’ (e.g., [10, 11, 32]), others prefer ‘identifications’ (e.g., [7, 12]), while still others refer to them as “just-is”-statements’ (e.g., [34]). Nothing philosophically significant turns on which label we select—so long as the targeted reading itself is clear. For the purposes of this paper, I will refer to these sentences as ‘generalized identities’. My use of the term ‘analysis’ differs from ‘generalized identity’ only in that it is slightly more expansive: it captures both generalized identities and

sentences with different syntactic structures that resemble generalized identities in philosophically significant respects.

It is often possible to express analyses using verbs (and verb phrases) in various forms. Examples of these sorts of sentences include:

6. To know that p is to have a justified true belief that p .
7. To die is to cease to live.
8. Shrinking is decreasing in size.
9. To resemble is to be similar to.

Perhaps an exhaustive discussion of analysis ought to address these types of sentences as well. I presently have very little to say about them and largely disregard them here.

However, I maintain that sentences of the form ‘For a to be F is for a to be G ’ belong to the same family as generalized identities—primarily because of their use in expressing analyses of 0-ary predicates. ‘To be F is to be G ’ may, and often does, express analyses of predicates of various adicities. Both ‘To be a vixen is to be a female fox’ and ‘To be adjacent to is to be next to’ are grammatically correct, although ‘be a vixen’ is a monadic predicate, while ‘be adjacent to’ is dyadic. But sentences of this form are typically ungrammatical when applied to 0-ary predicates (i.e., predicates that apply to no objects).¹ For example, ‘To be John is a bachelor is to be John is an unmarried male’ is ungrammatical. However, ‘For John to be a bachelor is for John to be an unmarried male’ is perfectly grammatical, and conveys the information expected of an analysis.

Under the target reading, the ‘is’ of generalized identity shares many logical and modal features with the ‘is’ of identity. It is reflexive, symmetric, and transitive, and if a sentence of this form is true, then it is necessarily true, and necessary that all and only F s are G s.² An adequate account of generalized identity ought, minimally, to explain the presence of these features.

Some dispute that ‘To be F is to be G ’ is reflexive and symmetric, maintaining that while ‘To be a father is to be a male parent’ is true, ‘To be a male parent is to be a father’ is false, and that ‘To be a father is to be a father’ is trivially false, rather than trivially true.³ Such philosophers hold that generalized identities form a strict partial ordering over predicates, perhaps maintaining that if ‘To be F is to be G ’ is true, then G is somehow more basic or fundamental than F is.

I confess that I once had such predilections myself. However, I have come to endorse a reading of ‘To be F is to be G ’ that strongly resembles an identity. There is a perfectly intelligible reading of ‘To be F is to be F ’ on which the sentence is manifestly true. After

¹ I make the standard assumption that sentences are 0-adic predicates. So, just as ‘smaller than’ is a dyadic predicate and ‘smaller than Jones’ is a monadic predicate, ‘Smith is smaller than Jones’ is a 0-adic predicate.

² The necessity of coextensionality relates to identity via (the predicative analogue of) Leibniz’s Law. At least when opaque predicates are set aside, the identity of Hesperus and Phosphorus guarantees that Hesperus and Phosphorus bear the same properties—that ‘Hesperus is a planet’ holds just in case ‘Phosphorus is a planet’ holds as well. As it concerns predicative identity, Leibniz’s Law ensures that F bears the higher-order property *contains a within its extension* just in case G bears that property as well. And since Leibniz’s Law is presumably not only true but necessarily true, every possible situation in which a is in the extension of F is one in which a is within the extension of G . Presumably, this strategy only succeeds if the predicate ‘contains a within its extension’ is not opaque. For recent discussions regarding the logic of opaque predicates in this context, see [4, 5, 7].

³ For one such argument, see [8].

all, what else could *F* possibly be?⁴ Of course, uttering this sort of sentence is typically pragmatically infelicitous. When someone inquires, ‘What is it to be a bachelor?’ the response ‘To be a bachelor is to be a bachelor’ is likely to be unhelpful—in much the same way that the responding to ‘Who is Bob?’ with ‘Bob is Bob’ is uninformative.

In addition, interest in one reading need not preclude the import of another. Although I am primarily concerned with the symmetric and reflexive reading of ‘To be *F* is to be *G*’, on the present approach it is possible to define an asymmetric and irreflexive notion of analysis in terms of the symmetric and reflexive one. I outline an account of such a reading in some concluding remarks and ask readers primarily concerned with this alternate reading to bear with me until then.

Once the reflexive and symmetric reading is embraced, it is natural to treat the ‘is’ of ‘To be *F* is to be *G*’ literally as the ‘is’ of identity (i.e., to maintain that a sentence of this form is true just in case the property of *being F* is identical to the property of *being G*). Notably, however, philosophers have objected to this suggestion for various reasons.⁵ First, there is linguistic evidence that ‘to be *F*’ is not synonymous with ‘to be the property of *being F*’ in some contexts. ‘I hope to be a skilled philosopher’ is perfectly true, but ‘I hope to be the property of *being a skilled philosopher*’ is presumably false; I do not hope to be a property. If so, ‘To be *F* is to be *G*’ may not be synonymous with ‘To be the property of *being F* is to be the property of *being G*’. Second, some desire an account of generalized identity to be compatible with nominalism—the denial that abstract objects such as properties and relations exist. A nominalist may grant that to be a triangle is to be a polygon with three angles, but would presumably deny that the property of *being a triangle* is identical to the property of *being a polygon with three angles* on the grounds that she denies the existence of properties. Minimally, some suggest, such a nominalist’s position ought not be obviously incoherent. By denying that ‘To be *F* is to be *G*’ is synonymous with ‘To be the property of *being F* is to be the property of *being G*’, such a nominalist’s position can be accommodated. Third, the nominalized version of generalized identity arguably possesses the resources to avoid Russell-style paradox. It may be that to be a non-self-instantiator is to not instantiate oneself, but there are reasons to deny that to be the property of *being a non-self-instantiator* is to be the property of *not instantiating oneself*, on the grounds that the property of *being a non-self-instantiator* gives rise to paradox. For these reasons, those like Dorr and Correia deny that ‘To be *F* is to be *G*’ is synonymous with ‘To be the property of *being F* is to be the property of *being G*’.

I find these arguments deeply unpersuasive. I am unprepared to settle central metaphysical debates based on the infelicity of ‘I hope to be the property of *being an accomplished philosopher*’—it may be that ‘to be *F*’ functions differently in that type of construction than it does in ‘To be *F* is to be *G*’. Moreover, those who interpret a theory of generalized identity as an account of property identity need not maintain that ‘to be *F*’ is synonymous with ‘to be the property of *being F*’. After all, the claim of synonymy need not follow from the claim of property identity. And while some accounts

⁴ Likewise, while there are contexts in which ‘To be a father is to be a male parent’ seems appropriate, there are also contexts in which ‘To be a male parent is to be a father’ seems appropriate. See [12] for a discussion of this point. And even those who primarily address an irreflexive reading of ‘To be *F* is to be *G*’ often countenance a reflexive reading as well (e.g., [35]).

⁵ So far as I can tell, [12] and [10] provide precisely the same arguments.

of generalized identity rule out forms of nominalism, I do not find this particularly undesirable. It would be naïve to expect a theory of analysis to stay neutral on all other philosophical debates. An account may regularly, perhaps even systematically, force our hand; so long as our hand is forced in the correct way, this seems no cost to me. Furthermore, while the reified approach entails that nominalism is incorrect, it does not entail that nominalism is *obviously* incorrect, because even if the reified approach is true, it need not be obviously true. And so, nominalists may be making a substantive, rather than a silly, mistake. And while the threat of paradox is real, it is a threat that philosophers of all stripes must face.

For the purposes of this paper, I treat ‘To be *F* is to be *G*’ in the Platonist way: in a manner that seeks identity conditions for properties, propositions, and relations. It is my hope that the target phenomenon is, by this point, sufficiently clear. I now turn to developments in truth-maker semantics, which I employ within this account.

§3. Truth-maker semantics. Truth-maker semantics rests upon the conviction that the meanings of sentences are determined by the aspects of the world that are responsible for their truth-values. When stated so generally, this hardly seems controversial. Who would doubt that the meaning of ‘Grass is green’ depends on that which is responsible for the truth of ‘Grass is green’? And, to be fair, many accounts claim conformity to this mantra—at least when it is understood broadly enough. The conviction that sentences’ meanings are associated with their truth conditions has been around at least since [36] and arguably since [26].

What is unique to truth-maker semantics is its commitment to *exact* truth-makers. The aspects of the world that constitute the meaning of a sentence *s* are those aspects which *exactly* make *s* true or *exactly* make *s* false; that is, the aspects of the world that determine *s*’s meaning are entirely relevant to *s*. By design, truth-maker semantics eliminates aspects of sentences’ meanings unrelated to their truth-values and is capable of attributing distinct meanings to necessarily equivalent sentences.⁶

Integral to the truth-maker approach is the idea of a *state*—a fact-like entity that typically has a more restricted scope than a possible world. For example, there is a state of whales being mammals, a state of it raining outside, and a state of 2 being a number. Additionally, there is a state of their being 12 planets in the solar system and a state of 2 being an odd number. So, states are not restricted to those that actually obtain, nor even to those that possibly could obtain. ‘State’ is intended as a term of art and is as metaphysically neutral as possible. I make no assumptions about what ontological kind of thing a state is, or about whether states figure in the fundamental building blocks of reality. One of the only metaphysical assumptions I make is that states are capable of mereological composition. So, the state of a ball being both red and round may be the composite of the state of the ball being red with the state of the ball being round.

Let a *state space* be an ordered pair $\langle S, \sqsubseteq \rangle$, where *S* is a set of states and \sqsubseteq is a binary relation on *S*. \sqsubseteq is intended to be interpreted as the relation of parthood, such that ‘ $s \sqsubseteq s'$ ’ asserts that state *s* is a part of state *s'*. I make the standard assumption that parthood is a partial ordering—that is, that \sqsubseteq satisfies the following criteria:

⁶ To a very large extent, this paper relies on truth-maker semantics as it was developed by [19–21, 23, 24]. For precursors to this semantics, see [28, 37].

- a) Reflexivity: For any state $s \in S$, $s \sqsubseteq s$;
- b) Antisymmetry: For any states $s, s' \in S$, if $s \sqsubseteq s'$ and $s' \sqsubseteq s$, then $s = s'$;
- c) Transitivity: For any states $s, s', s'' \in S$, if $s \sqsubseteq s'$ and $s' \sqsubseteq s''$, then $s \sqsubseteq s''$.

Many state spaces, as defined above, are uninteresting. For example, there are state spaces in which no mereological composition occurs. In these, the extension of \sqsubseteq is restricted to reflexivity; every state is a part of itself, and no state is a part of any other. For the purposes of this paper, let us confine our attention to state spaces that allow for arbitrary fusion, which I call ‘complete state spaces’. For the most part, this can be achieved by stipulating that any two states within S have a fusion within S . However, this approach fails for infinitely large state spaces. A restriction appropriate for infinitely large state spaces requires a few more definitions.

I denote the fusion of states $T \subseteq S$ as $\bigsqcup T$. Let an *upper bound* of $T \subseteq S$ be a state s such that, for all states $t \in T$, $t \sqsubseteq s$. That is to say, an upper bound of a subset of S is a state that contains—as a part—every state within that subset. Let a *least upper bound* of $T \subseteq S$ be a state s such that s is an upper bound of T , and for all upper bounds s' of T , $s \sqsubseteq s'$. Intuitively, we can think of the least upper bound of T as being the smallest upper bound of T —one that is a part of all upper bounds of T . Provably, if there is a least upper bound of T , then there is a unique least upper bound of T . For a state space $\langle S, \sqsubseteq \rangle$, select an arbitrary $T \subseteq S$. Suppose, for reductio, that T has two least upper bounds— $\bigsqcup T^1$ and $\bigsqcup T^2$. From the definition of ‘least upper bound’ we have that $\bigsqcup T^1 \sqsubseteq \bigsqcup T^2$ and $\bigsqcup T^2 \sqsubseteq \bigsqcup T^1$. Given antisymmetry, this entails that $\bigsqcup T^1 = \bigsqcup T^2$. Let a state space $\langle S, \sqsubseteq \rangle$ be *complete* just in case every subset $T \subseteq S$ has a fusion within S . Here, I consider only complete state spaces. Furthermore, I identify the fusion of states with the least upper bound of the set containing all and only those states.

On some approaches, the end of a semantics is to determine whether statements are true or false in a given world. However, on the current approach, the aim is to determine what precisely it is within a world that is responsible for the truth-values of sentences. Thus, verifiers must be relevant to the sentences that they render true, and they must be entirely relevant. States are the candidate verifiers and falsifiers. So, the state of the Republicans controlling the Senate verifies ‘The Republicans control the Senate’ and falsifies ‘The Republicans do not control the Senate’.

I do not assume that each sentence has a unique verifier and a unique falsifier. Potentially, sentences could have many. The sentence ‘Either Dante wrote *The Divine Comedy* or Aristotle tutored Alexander the Great’ presumably has (at least) two verifiers: the state of Dante having written *The Divine Comedy* and the state of Aristotle having tutored Alexander the Great. However, it is important for a sentence’s verifiers to guarantee its truth and for its falsifiers to guarantee its falsehood. So, the state of roses being red does not verify ‘Roses are red and violets are blue’, despite its relevance to the conjunction. Furthermore, verifiers and falsifiers do not contain extraneous information. The state of 2 being both even and prime does not verify ‘The number 2 is even’, despite the fact that it guarantees the sentence’s truth.⁷

⁷ This differentiates truth-maker semantics from the approach outlined by [6] and [30] who employ a notion of truth-making such that verifiers need not be wholly relevant to the sentences that they render true.

Nothing in the definition of a state space distinguishes those states that are possible from those that are impossible. Given the modal implications of generalized identities, it is valuable to be able to make such a distinction. Let a *modalized state space* be an ordered triple $\langle S, S^\circ, \sqsubseteq \rangle$ such that $\langle S, \sqsubseteq \rangle$ is a complete state space, and S° is a nonempty subset of S consisting of those states that are possible. I assume that S° is closed under parthood; every possible state s is such that all of its parts are possible states ($s \in S^\circ \wedge s' \sqsubseteq s \rightarrow s' \in S^\circ$).

Modalized state spaces have the resources to define compatibility and incompatibility. Let us say that states s and s' are *compatible* just in case their fusion is a member of S° , and are *incompatible* if it is not. It is also possible to employ modalized state spaces to define a notion of a possible world. A possible world is a maximally consistent modalized state space.⁸ So, if something holds in all modalized state spaces, it thereby holds in all possible worlds. Several other requirements might be imposed upon verification and falsification once states are assigned to sentences. We might, for example, insist that no verifier of a sentence is compatible with a falsifier of that sentence and that every possible state is compatible with either a verifier or a falsifier of every sentence (which corresponds to *exclusivity* and *exhaustivity* conditions of verification, respectively).

For obvious reasons, propositional languages are of limited relevance to claims of the form ‘To be F is to be G ’. Let us restrict our attention to a language capable of expressing predication. Let I be the (potentially infinite) set of all objects.⁹ Let a language L contain a unique name for each object such that \underline{i}_1 denotes i_1 , \underline{i}_2 denotes i_2 , etc. Additionally, let L contain infinitely many predicates P_1, P_2, \dots of fixed adicity; the operators \neg, \wedge , and \vee ; infinitely many variables, and the quantifiers \exists, \forall —all defined in the standard way.

Let a model M be an ordered quadruple $\langle S, \sqsubseteq, I, |\cdot| \rangle$ such that $\langle S, \sqsubseteq \rangle$ is a complete state space, I is the set of individuals, and $|\cdot|$ is a valuation function that takes each n -adic predicate P and each n -tuple of objects $\langle i_1, i_2, \dots, i_n \rangle$ to an ordered pair $\langle V, F \rangle$ where both V and F are subsets of S , with the intended interpretation that V is the set of $P(\underline{i}_1, \underline{i}_2, \dots, \underline{i}_n)$ ’s verifiers, and F is the set of its falsifiers. The semantics is given inductively:

- i)⁺ $s \Vdash P(\underline{i}_1, \underline{i}_2, \dots, \underline{i}_n)$ iff $s \in |(P(i_1, i_2, \dots, i_n))|^V$;
- i)⁻ $s \Vdash P(\underline{i}_1, \underline{i}_2, \dots, \underline{i}_n)$ iff $s \in |(P(i_1, i_2, \dots, i_n))|^F$;
- ii)⁺ $s \Vdash \neg A$ iff $s \not\Vdash A$;
- ii)⁻ $s \Vdash \neg A$ iff $s \Vdash A$;
- iii)⁺ $s \Vdash A \wedge B$ iff for some states t and u , $t \Vdash A$ and $u \Vdash B$ and $s = t \sqcup u$;

⁸ For a given modalized state space $\langle S, S^\circ, \sqsubseteq \rangle$ let a *subspace* be a modalized state space $\langle S', S'^\circ, \sqsubseteq' \rangle$ such that $S' \subseteq S, S^\circ \subseteq S'^\circ$, and \sqsubseteq' is a conservative extension of \sqsubseteq defined over S' . A subspace $\langle S', S'^\circ, \sqsubseteq' \rangle$ of $\langle S, S^\circ, \sqsubseteq \rangle$ is a *possible world* just in case $S^\circ = S'^\circ$, $\sqcup S' \in S^\circ$, and every set S'' such that $S' \subseteq S'' \subseteq S$ is such that $\sqcup S'' \notin S^\circ$.

⁹ Orthodox set theory denies that there is a set of all objects, for such a set would presumably contain itself (leading quickly to the Russell Paradox). For this reason, the use of ‘all’ in ‘set of all objects’ is, presumably, restricted. Indeed, it is for similar reasons that some maintain that *all* quantifiers are restricted—see [13, 18, 31, 33]—though it is presumably challenging to express the claim that all quantifiers are restricted in a satisfactory way. I set the issue of *how* the ‘all’ ought to be restricted aside (though presumably it will at least be unable to quantify over that very set). I also note that nonstandard set theories allow for a universal set.

- iii)⁻ $s \Vdash A \wedge B$ iff either $s \Vdash A$ or $s \Vdash B$;
- iv)⁺ $s \Vdash A \vee B$ iff either $s \Vdash A$ or $s \Vdash B$;
- iv)⁻ $s \Vdash A \vee B$ iff for some states t and u , $t \Vdash A$ and $u \Vdash B$ and $s = t \sqcup u$.

It is my hope that readers find this semantics to be intuitive. Nevertheless, it has some surprising results. One of the most unexpected is that a verifier of $A \wedge A$ need not verify A . Every fusion of two distinct verifiers of A is a verifier of the conjunction $A \wedge A$, but these fusions need not themselves verify A . I do not find this result particularly problematic. Indeed, as I discuss in Section 4, it is directly responsible for desirable logical attributes of analysis. However, those who insist that a verifier of A ought to verify $A \wedge A$ can readily modify the semantics by requiring that verifiers are closed under fusion; the fusion of any two verifiers of A is itself a verifier of A . Clauses iii)⁻ and iv)⁺ thus become:

- iii)^{*-} $s \Vdash A \wedge B$ iff $s \Vdash A$ or $s \Vdash B$ or $s \Vdash A \vee B$;
- iv)^{*+} $s \Vdash A \vee B$ iff $s \Vdash A$ or $s \Vdash B$ or $s \Vdash A \wedge B$.

There are at least two ways to extend the semantics to clauses with quantifiers. One utilizes generic objects, such that verifiers of universal statements are generic states (rather than states about particular objects). A verifier for ‘All numbers are either even or odd’ is the state of a generic number being either even or odd. The semantics I prefer is instantial: verifiers of universal and existential statements are states concerning their instances. The introduction of predicates is currently more significant than the introduction of quantifiers is, so I opt for the most easily intelligible approach to quantification. The instantial approach treats the semantics of universal statements like large conjunctions. The meaning of the claim ‘Everything is F ’ is treated similarly to the conjunction of the claims ‘ $F(i)$ ’, for all i . More formally, we have:

- v)⁺ $s \Vdash \forall x A(x)$ iff there is a function f from I into S such that $f(i) \Vdash A(i)$ for each $i \in I$, and $s = \bigsqcup \{f(i) : i \in I\}$;
- v)⁻ $s \Vdash \forall x A(x)$ iff for some $i \in I$, $s \Vdash A(i)$;
- vi)⁺ $s \Vdash \exists x A(x)$ iff for some $i \in I$, $s \Vdash A(i)$;
- vi)⁻ $s \Vdash \exists x A(x)$ iff there is a function f from I into S such that $f(i) \Vdash A(i)$ for each $i \in I$, and $s = \bigsqcup \{f(i) : i \in I\}$.¹⁰

¹⁰ One further complication arises if we assume that a verifier ought to necessitate the sentence it verifies, and allow for a variable domain of objects that could exist. Suppose, for example, that all cats have tails. On the instantial approach, a verifier of ‘All cats have tails’ is the fusion of a verifier of the claim that each cat has a tail—that is, that Fluffy is a cat with a tail, Scamp is a cat with a tail, and so on. This fusion could obtain and yet ‘All cats have tails’ be false in a possible situation in which there are more cats than actually exist—one of which lacks a tail. The obvious solution to this problem is to incorporate a totality state—a verifier of a universal statement is the fusion of verifiers of all instances with a verifier of the claim that those are all of the objects that exist. A similar solution applies to cases where objects (which actually exist) would not have existed. Plausibly, a verifier of ‘All objects exist’ is the fusion of the state of a existing with the state of b existing, and so on, with the totality state of a, b, \dots , being all of the objects. If a were not to exist, then this totality state would not obtain. However, there would remain a verifier of ‘All objects exist’ because the totality state of b, c, \dots , being all of the objects would then obtain, and the fusion of this state with the state of b existing, with the state of c existing, and so on, would verify ‘All objects exist’. For the development of this approach, see [22]. Interestingly, once the notion of a totality state is employed, it is possible to define the distinction between inclusive and

Exact equivalence can thus be defined as follows: sentences A and B are *exactly equivalent* just in case their verifiers and falsifiers are identical. Many logically equivalent sentences (from a classical perspective) are not exactly equivalent. For example, $A \wedge \neg A$ need not be exactly equivalent to $B \wedge \neg B$. Verifiers of the first sentence are fusions of a verifier of A with a falsifier of A , while verifiers of the second sentence are fusions of a verifier of B with a falsifier of B . And, as mentioned above, A is not even exactly equivalent to $A \wedge A$ unless appropriate modifications are made. Nevertheless, exact equivalence differs from syntactic identity. $A \wedge B$ is exactly equivalent to $B \wedge A$, and A is exactly equivalent to $\neg\neg A$.¹¹

§4. The semantic foundations of philosophical analysis. Truth-maker semantics provides the resources to account for ‘To be F is to be G ’ perspicuously: to elucidate what I meant when I claimed that these sentences hold just in case that which makes something F is that which makes something G . A sentence of this form is true just in case, for any name i , ‘ $F(i)$ ’ is exactly equivalent to ‘ $G(i)$ ’.¹² ‘To be a person is to be bound by the categorical imperative’ holds just in case the verifiers and falsifiers of the claim that someone is a person are identical to the verifiers and falsifiers of the claim that she is bound by the categorical imperative, and ‘To be morally right is to maximize utility’ holds just in case the verifiers and falsifiers of the claim that an act is morally right are identical to the verifiers and falsifiers of the claim that it maximizes utility.¹³

It may seem odd, given that I am ostensibly concerned with what makes it the case that an object is F , that I am sensitive to states that make it the case that something *isn't* F —that is, that this account is sensitive to falsifiers as well as to verifiers.¹⁴ On some semantic approaches, the inclusion of falsifiers is gratuitous. If the meaning of ‘ Fa ’ is given by the possible worlds in which a is F , this collection also determines the worlds in which a is not F . I suspect this gratuity is responsible for the apparent oddity of mentioning falsifiers explicitly. But on the truth-maker approach this is not the case.

exclusive quantification. For if we insist that there be a *nonempty* fusion of instances with totality states, the resulting quantifiers will be exclusive, but if we allow for empty fusions of instances with totality states, the resulting quantifiers will be inclusive. In the case of an empty domain with inclusive quantification, a verifier of the claim that everything is F will be the state of there being nothing that exists.

¹¹ This itself is a desirable outcome; it would be something of a bizarre coincidence if the granularity of properties, propositions, and relations perfectly aligned with the granularity of our syntax.

¹² It is straightforward to generalize this account to apply to predicates of any fixed adicity. For any pair of n -adic predicates F and G , the sentence ‘To be F is to be G ’ is true just in case, for all collections of n names, $i_1—i_n$, ‘ $F(i_1, i_2, \dots, i_n)$ ’ is exactly equivalent to ‘ $G(i_1, i_2, \dots, i_n)$ ’.

¹³ Perhaps some worry that this account is specified in terms of names. Arguably, this requires every object within the domain to have a name. This may well require the language to be equipped with uncountably many names (for example, if the account is applicable to all of the real numbers). This is dubious if languages are to be understood as the sorts of things that ordinary speakers understand. This worry can be avoided by shifting the account to one that appeals directly to objects, rather than to names of objects. That is, the semantics for atomic sentences is altered so that it applies to predicates and objects, rather than predicates and names. The valuation function may then be extended inductively so that it applies to formulas with free variables (not merely atomic predicates), and ‘To be ϕ is to be ψ ’ is true iff $\forall a \in I, |\phi(a)|^V = |\psi(a)|^V$ and $|\phi(a)|^F = |\psi(a)|^F$.

¹⁴ The first truth-maker approach sensitive to falsification occurs in [21].

Sentences may be alike in what makes them true but differ in what makes them false. It seems to me that there is nothing to be gained from the exclusion of falsifiers, and a great deal is lost. For, as we shall see, including falsifiers allows one to infer from the fact that the falsifiers of ‘Sarah is a bachelor’ differ from the falsifiers of ‘Sarah is a vixen’ that ‘To be a bachelor is to be a vixen’ is false, where an approach sensitive only to verification cannot justify this inference.¹⁵

It is intended for this account to have a wide scope. It applies to predicates of any fixed adicity (i.e., with any number of free variables). This includes not only unary predicates (such as ‘is tall’) and binary predicates (such as ‘is next to’) but also the limiting case of 0-ary predicates. ‘For John to be a brother is for John to be a male sibling’ holds just in case the states that verify and falsify ‘John is a brother’ are identical (respectively) to the states that verify and falsify ‘John is a male sibling’. However, it is worth emphasizing that what has been provided is merely a *first-order* semantics. It is not an account of higher-order generalized identity; it does not, for example, provide conditions in which ‘To be a property is to be a universal that applies to objects’ is true. This is in contrast to other approaches in the literature—many of which apply to infinitely higher-order languages as well. This limitation is largely due to the current state of truth-maker semantics. It has been developed as a semantics for first-order languages, but it is not at all obvious how it is to be expanded to higher-order languages. If (and when) a truth-maker semantics for higher-order languages is developed, it is my hope that this account of generalized identity may be expanded in the obvious way.

Some might suspect that sentences of the form ‘To be F is to be G ’ are true just in case ‘ $\forall xFx$ ’ is exactly equivalent to ‘ $\forall xGx$ ’. This is a quite natural approach—it is also false. ‘ $\forall xFx$ ’ is exactly equivalent to ‘ $\forall xGx$ ’ just in case the sentences’ verifiers and falsifiers are identical. On the current approach, universal statements are treated instantially. That is to say, the claim that everything is F is treated like the conjunction of the claims that each individual is F . Suppose that there were two objects i_1 and i_2 such that a verifier for ‘ $F(i_1)$ ’ was a verifier for ‘ $G(i_2)$ ’, and a verifier for ‘ $F(i_2)$ ’ was a verifier for ‘ $G(i_1)$ ’. In this case, the conjunction ‘ $F(i_1) \wedge F(i_2)$ ’ has an identical verifier to ‘ $G(i_1) \wedge G(i_2)$ ’—the verifier for each is the fusion of two identical states. Nevertheless, it would be absurd to take this to lend support to the truth of ‘To be F is to be G ’. After all, a state that renders ‘ $F(i_1)$ ’ true is a state that renders ‘ $G(i_2)$ ’, rather than ‘ $G(i_1)$ ’, true. It is essential to preserve the verifiers and falsifiers for each instance—and not merely for the fusion of their instances.¹⁶ Or consider predicates F

¹⁵ One may, in this case, conclude that ‘To be a bachelor is to be a vixen’ is false on the grounds that ‘Sarah is a bachelor’ has different verifiers from ‘Sarah is a vixen’ (although it may be that, in the case in which Sarah is a female human, neither of these states obtains). However, it seems that one need not appeal to verifiers in order to conclude that ‘To be a bachelor is to be a vixen’ is false; someone armed merely with the knowledge that ‘Sarah is a bachelor’ differs in falsifiers from ‘Sarah is a vixen’ has the resources to conclude that ‘To be a bachelor is to be a vixen’ is false. This is something that the bilateral approach (which appeals to both verifiers and falsifiers) accommodates and the unilateral approach (which appeals only to verifiers) does not.

¹⁶ This problem may be peculiar to the instancial approach to quantification. On the generic approach—according to which verifiers of ‘ $\forall xFx$ ’ are states that verify that a generic object is F —there is no worry about swapping verifiers of instances. My thanks to Kit Fine for pressing me on this point. However, once the generic approach is adopted, the ability to define the distinction between inclusive and exclusive quantification in the manner mentioned in

and G where G is the predicate ‘is such that everything is F ’; for the sake of concreteness, let F be ‘is tall’ and G be ‘is such that everything is tall’. It may be that ‘Everything is tall’ is exactly equivalent to ‘Everything is such that everything is tall’, but it does not follow from this that to be tall is to be such that everything is tall.¹⁷

This account accommodates the aforementioned logical features of generalized identity, primarily because exact equivalence is an equivalence relation. For any F and any i , ‘ $F(i)$ ’ is exactly equivalent to ‘ $F(i)$ ’, so sentences of the form ‘To be F is to be F ’ are universally true. Exact equivalence is symmetric; if ‘ $F(i)$ ’ is exactly equivalent to ‘ $G(i)$ ’, then ‘ $G(i)$ ’ is exactly equivalent to ‘ $F(i)$ ’. So, if ‘To be F is to be G ’ is true, then ‘To be G is to be F ’ is true. Exact equivalence is also transitive. If ‘ $F(i)$ ’ is exactly equivalent to ‘ $G(i)$ ’ and ‘ $G(i)$ ’ is exactly equivalent to ‘ $H(i)$ ’, it follows that ‘ $F(i)$ ’ is exactly equivalent to ‘ $H(i)$ ’. So, if ‘To be F is to be G ’ is true and ‘To be G is to be H ’ is true, then ‘To be F is to be H ’ is also true. Generalized identities are thus reflexive, symmetric, and transitive.

Of course, with a semantics in place we can establish more than that exact equivalence is an equivalence relation: it determines a logic of analysis.¹⁸

The present account entails that sentences of the following forms (which correspond to axioms in a logic of exact equivalence) are universally true and that inferences of the following forms (which correspond to inferences in a logic of exact equivalence) are universally valid:

Axioms

- | | |
|---|---|
| 1. To be $\neg\neg F$ is to be F . | 5. To be $F \vee (G \vee H)$ is to be $(F \vee G) \vee H$. |
| 2. To be $F \wedge G$ is to be $G \wedge F$. | 6. To be $\neg(F \wedge G)$ is to be $\neg F \vee \neg G$. |
| 3. To be $F \vee G$ is to be $G \vee F$. | 7. To be $\neg(F \vee G)$ is to be $\neg F \wedge \neg G$. |
| 4. To be $F \wedge (G \wedge H)$ is to be $(F \wedge G) \wedge H$. | |

Inferences

8. To be F is to be G /To be G is to be F .
9. To be F is to be G /To be $\neg F$ is to be $\neg G$.
10. To be F is to be G , To be G is to be H /To be F is to be H .
11. To be F is to be G /To be $F \wedge H$ is to be $G \wedge H$.
12. To be F is to be G /To be $F \vee H$ is to be $G \vee H$.

This account relies upon a connection between sentential and predicative operators. While truth-maker semantics has been developed with an eye to sentential connectives (perhaps providing verifiers for ‘John is unmarried and John is male’, for example),

footnote 10 is lost. Additionally, the instantial approach has the advantages of comparative parsimony (in that it need not posit generic objects) and relative intelligibility when compared to generic verification.

¹⁷ My thanks to an anonymous reviewer for this example.

¹⁸ For a discussion of the logic of exact equivalence, see [25]. I do not here establish whether the logic is complete on the present semantics. [21], however, establishes the completeness of a closely related system. He suggests that close variants of his system are also complete, claiming, “We shall prove completeness of the system by way of disjunctive normal forms. This is perhaps the simplest way to prove completeness within the present context, though it does not readily apply to various extensions of the system, for which other methods of proof may be used” (p. 15). Because neither Fine nor I provide a completeness theorem for this precise logic, I make no claim of completeness here.

accounts of generalized identity are primarily concerned with predicates (perhaps maintaining that to be a bachelor is to be an unmarried male). It should be no surprise that there is a close connection between the two. Sentences are often conceived of as the limiting case of predicates (i.e., predicates that take 0 objects within their sentences), so—minimally—accounts of predicates ought to apply to sentences as well.

Formalizations of logically complex predicates often employ λ -notation. The predicate $F \wedge G$, on this approach, might be represented as $\lambda x.(Fx \wedge Gx)$. The intended connection to truth-maker semantics is that $\lambda x.(Fx \wedge Gx)a$ has the same verifiers as $Fa \wedge Ga$. This follows from two standard assumptions about propositional identity: that propositional identity is preserved under β -conversion and that Leibniz's Law holds. For, on the assumption that identity is preserved under β -conversion, $\lambda x.(Fx \wedge Gx)a$ simply is $Fa \wedge Ga$, and Leibniz's Law entails that identical propositions have identical verifiers.¹⁹

I believe this account to be true for several reasons. Two relatively banal ones are that it is intuitive and that it satisfies our theoretical demands with minimal theoretical costs. But perhaps the most intriguing reason I believe that this is so, and a guiding motivation behind the current approach, concerns the analogy between real and nominal definition: the analogy between the definitions of things themselves and the definitions of terms (for example, the analogy between the definition of the good and the definition of 'the good'). On one conception, nominal definitions are a special kind of real definition—they are the real definitions of linguistic objects.²⁰ Even those who do not conceive of nominal definition in this way often countenance a close connection between it and real definition. It is common practice for the notion of real definition to be introduced as the worldly analogue to (or perhaps the metaphysical counterpart of) nominal definition. If the two are so closely related, then the philosophically salient aspects of one ought to resemble the philosophically salient aspects of the other. Of course, this analogy determines little about real definition without an account of nominal definition—a theory of how words, phrases, sentences, and the like are defined. One popular conception of nominal definition is semantic essentialism: the view that terms are defined by their semantic content.²¹ The definition of a word is given by what the word means, and the definition of a sentence is given by what it is the sentence means.

If we are to take both semantic essentialism and the analogy between real and nominal definition seriously (and I strongly believe that we should), then the logical attributes of philosophical analysis ought to mirror the logical attributes of meaning (or, some might say, the logical attributes of analyticity). After all, on this view, real definitions resemble the way in which terms are defined, and terms are defined by their

¹⁹ The assumption that propositional identity is preserved under β -conversion is orthodox (for a defense, see [12], for those who assume it to be true, see, e.g., [7, 27]). However, this principle is incompatible with the claim that Fa grounds $\lambda x.(Fx)a$ (at least on the assumption that grounding is irreflexive). I suspect that those concerned with this type of grounding principle would find the discussion of real definition at the conclusion of this paper more appealing than the account that resembles an identity. Furthermore, I note that there have recently been substantial discussion regarding systems that abandon Leibniz's Law in this context—see, again [7], as well as [4, 5].

²⁰ On another conception, nominal definitions can be thought of as putative real definitions that are defective, or in some way of inferior status. I set this use of 'nominal definition' aside.

²¹ See, e.g., [17].

meanings. To put it only slightly differently, if the logic of real and nominal definition functioned radically differently, I would conclude that they are disanalogous in the relevant sense. As it turns out, we already possess a logic of meaning (see [1–3]). Angell argues that this logic is nonclassical. Although p classically entails $p \vee q$, there is no reason to suspect that the meaning of p contains the meaning of $p \vee q$. For this reason, disjunction introduction fails on this logic. Angell's system has proven valuable in constructing a theory of partial truth. There is a sense in which 'Roses are red and violets are orange' is merely partially true. As it turns out, it is extremely difficult to construct a theory of partial truth with the resources of classical logic. One might suspect that a sentence is merely partially true just in case it entails something that is true and entails something else that is false. But this is too permissive: every strictly false sentence would count as being merely partially true. After all, every false sentence entails infinitely many logical truths. One might, instead, suggest that a sentence is merely partially true just in case it entails a *nontrivial* truth and a *nontrivial* falsehood, but this sacrifices too much. The sentence 'Grass is red and $1 + 1 = 2$ ' would not count as being merely partially true, because the relevant true conjunct (that $1 + 1 = 2$) is trivial. It is much more straightforward to construct a theory of merely partial truth on Angell's system. A sentence is merely partially true, on this account, just in case at least one of its analytic parts—a part of its meaning—is true and at least one of its analytic parts is false. Partially because disjunction introduction fails on this approach, this avoids the implication that every sentence is at least partially true. It is a virtue of my account that the logic of analysis aligns with Angell's system.

It is worth examining the plausibility of this system in detail. Each axiom and inferential rule is independently worthy of consideration, but I will primarily address three differences between the current approach and that advanced by [11], perhaps the most similar account available in the literature. In some ways, our approaches are incommensurable. While I provide a semantics that underlies the logic of analysis—one rooted in the mereological structures of states—Correia and Skiles do not. While they devote considerable discussion to metaphysical essence and ground, I do not. The point of overlap lies at the logic of generalized identity, and it is here that our views compete. Given that the target phenomenon is the same, it is unsurprising that our logics share several axioms and inferential rules. However, they come apart in important—and, I believe, quite telling—ways.

The first (and in some ways the least persuasive) difference between our systems concerns the distribution axioms. Correia and Skiles license 'To be $F \wedge (G \vee H)$ is to be $(F \wedge G) \vee (F \wedge H)$ ' but deny that 'To be $F \vee (G \wedge H)$ is to be $(F \vee G) \wedge (F \vee H)$ '. That is to say, on their system the distribution of conjunction over disjunction succeeds, while distribution of disjunction over conjunction fails. In contrast, I deny both distribution axioms. Their failure falls immediately out of the semantics I provide. One kind of verifier of ' $(Fa \vee Ga) \wedge (Fa \vee Ha)$ ' is the fusion of a verifier of ' Fa ' with a verifier of ' Ha '. This fusion need not verify ' $Fa \vee (Ga \wedge Ha)$ '. Similarly, one kind of falsifier of ' $(Fa \wedge Ga) \vee (Fa \wedge Ha)$ ' is the fusion of a falsifier of ' Fa ' with a falsifier of ' Ha '. This fusion need not falsify ' $Fa \wedge (Ga \vee Ha)$ '. So, neither distribution axiom is valid on my account.

I find the acceptance of only one distribution axiom bizarre. To be clear, while I claim both axioms are false, I do not find it bizarre to suggest that they are both true. What I find bizarre is accepting one while rejecting the other. Indeed, I can think of no

philosophical motivation for this result.²² Conjunction is symmetric with disjunction in logically important ways; the relation conjunction stands in to the True and the False is analogous to the relation disjunction stands in to the False and the True. And so, the two logical operators ought to distinguish predicates from one another in a parallel manner. This symmetry is broken on Correia and Skiles's account. A fundamental difference arises with the distribution axioms: while distributing conjunction over disjunction does not result in a different predicate, distributing disjunction over conjunction does.

I suspect that the inclusion of only one distribution axiom arises from their intended semantics. If a unilateral notion of exact equivalence were adopted—one that included verification and excluded falsification—only one distribution axiom would result. However, this is no defense unless a semantics that ignores falsification is independently motivated. Absent a reason to adopt a unilateral semantics, we remain without a reason to accept only one distribution axiom. The desire for a symmetric logic is unlikely to persuade those deeply wedded to an asymmetric account. At best, it might motivate uncertain philosophers waffling between available alternatives. But it remains a motivation for me, and is a harbinger of more serious problems to come.

A second (and perhaps more compelling) difference concerns an inference that I license and that Correia and Skiles do not—inference 9. If to be F is to be G , it follows that to be $\neg F$ is to be $\neg G$. For example, if to be a sister is to be a female sibling, then to be not a sister is to be not a female sibling, and if to be hydrogen is to be the chemical element with a single proton, then to be not hydrogen is to be not the chemical element with a single proton. While Correia and Skiles may (or may not) agree with these particular cases, they deny that the inference holds in its full generality.

Inference 9 also results from adopting a bilateral notion of exact equivalence (i.e., from insisting that sentences are exactly equivalent just in case they have both identical verifiers and identical falsifiers, rather than merely having identical verifiers).²³ To see

²² Notably, however, [9] provides putative counterexamples to only one distribution axiom (see also [29]). If there are counterexamples to one (but not both) distribution axiom, this might undercut the apparent oddity of accepting one and rejecting the other. Correia's case is framed in terms of fitting representation, rather than truth-making, but an analogous case can be straightforwardly derived. Let us suppose that Sam is both sad and ill. In this case, the state of Sam being both sad and ill is fittingly described (to use his terminology) by '(Sam is sad or Sam is bad) and (Sam is sad or Sam is ill)' but is not fittingly described by 'Sam is sad or (Sam is bad and Sam is ill)'. In my mind, this putative counterexample is most naturally viewed as a post-theoretic judgment, rather than an independently motivated consideration. This is evident from the fact that Correia employs the very logic he develops when motivating its result; it receives no independent justification—one that does not appeal to his very system. If this is permissible, it is straightforward to construct parallel cases for falsification. Let us suppose that Sam is neither bad nor sad. In this case, the state of Sam being neither bad nor sad is directly responsible for the falsity of '(Sam is bad and sad) or (Sam is bad and ill)' yet need not be directly responsible for the falsity of 'Sam is bad and (Sam is sad or ill)'. For although such a state guarantees the falsity of 'Sam is bad', it contains information not relevant to Sam's badness, and it likewise does not account for the falsity of 'Sam is sad or ill', as it has no information regarding Sam's illness. If the account I defend is correct, these are the sorts of counterexamples that are to be expected.

²³ Perhaps some suspect that there is no need to take a stand on negative generalized identities at all. It may seem metaphysically unimportant what it is to be $\neg F$, as the things that are not F do not form a uniform kind. It seems to me, however, that there are true negative generalized identities—though these sometimes involve quantifiers. It is plausible that for Alex to be an only child is for there not to exist someone who is a sibling of Alex (or, given

why this is so, consider the proposal that a generalized identity is true just in case, for any name i , ' $F(i)$ ' has identical verifiers to ' $G(i)$ ', and omitting the inclusion of falsifiers. Nothing in truth-maker semantics guarantees that sentences with identical verifiers have identical falsifiers. The negation operator swaps a sentence's verifiers for its falsifiers, so, in this case, ' $\neg F(i)$ ' could have distinct verifiers from ' $\neg G(i)$ ', and 'To be $\neg F$ is to be $\neg G$ ' would thereby be false. In contrast, my approach takes the false just as seriously as the true. If the states that falsify ' Fa ' differ from those that falsify ' Ga ', this undermines the claim that 'To be F is to be G '. And because generalized identities are equally sensitive to that which verifies that an object is F as they are to that which falsifies that an object is F , the inference from 'To be F is to be G ' to 'To be $\neg F$ is to be $\neg G$ ' is preserved.

An example may highlight the implausibility of denying inference 9. Suppose a precocious doctoral candidate were to announce that she had discovered that to be water is to be the chemical compound H_2O . "What an amazing finding!" her advisor might reasonably exclaim. "I just have one quick follow-up—what is it to be not water?" It would seem strange if such a student lacked the resources for the answer; it seems perfectly obvious that to be not water is to be not the chemical compound H_2O , and that this is immediately justified by the inference from the fact that to be water is to be the chemical compound H_2O .

If inference 9 were impermissible, I doubt she could answer her advisor's question at all. What resources could she employ to uncover what it is to be not water? Are there other scientific experiments than those she used to discover the chemical composition of water? If not, what inferential resources are available to her, if even the inference from 'To be F is to be G ' to 'To be $\neg F$ is to be $\neg G$ ' is denied? I suspect that those who deny inference 9 arrive at a limited (and implausible) form of skepticism: that there is no way for the student to uncover what it is to be not water. Minimally, the ball is in such philosophers' court: either show how the limited inferential resources can generate knowledge of what it is to be not water, or else accept the skepticism that results.

Notably, there are no counterexamples to inference 9 available in the literature. However, during the course of this paper's development various philosophers have presented putative counterexamples to me. I will mention two. It is natural to suggest that to be definitely red is to be red.²⁴ There may be borderline cases of redness—at times reasonable people disagree about whether an object is red. But there is consensus about the paradigmatic cases; we can be confident that definite redness is redness. However, even if this is true, 'To be not definitely red is to be not red' is arguably false, for things might be red that are not definitely red.

I maintain that if there is a reading of 'To be definitely red is to be red' on which the sentence is true, it is not the reading at issue. Recall that, on the target reading of 'To be F is to be G ', it is necessary that all *and only* F s are G s. So if there are possible instances of objects that are red but are not definitely red, then 'To be definitely red is

the symmetry of generalized identity, for there not to exist someone who is the sibling of Alex is for Alex to be an only child); it is plausible that for a number n to be prime is for there not to exist a number that n is evenly divisible by other than 1 and n (and, consequently, for there not to exist a number that n is evenly divisible by other than 1 and n is for n to be prime); it is plausible that to be immortal is to be such that there does not exist a time at which one dies.

²⁴ My thanks to Stephen Yablo for this example.

to be red' is false. In some ways this example resembles 'To be crimson is to be red', mentioned at the outset. For, just as being crimson is a way of being red (and, perhaps, being definitely red is a way of being red), there are other ways of being red than being crimson that render 'To be crimson is to be red' false on the target reading. Just so with 'To be definitely red is to be red'.

Another potential counterexample is the following: suppose, for the sake of simplicity, that there are only three colors—red, blue, and green—and let '*r*', '*b*', and '*g*' denote the states of a particular object *o* being red, blue, and green, respectively.²⁵ Arguably, the sentence '*o* is red' is verified by *r* and is falsified by *b* and *g* and by nothing else (and, arguably, the sentences '*o* is blue' and '*o* is green' are verified and falsified in the analogous way). Quite plausibly, in this case, to be not red is to be blue or green. But while the verifiers of '*o* is not red' may be the same as the verifiers of '*o* is blue or *o* is green', they differ in their falsifiers. The only falsifier of '*o* is not red' is *r*, but the only falsifiers of '*o* is blue or *o* is green' are *r*, $r \sqcup b$, $r \sqcup g$, and $g \sqcup b$ —that is, the state of *o* being red, the state of *o* being both red and blue, the state of *o* being both red and green, and the state of *o* being both green and blue.

This is an interesting case (one more compelling than others I have encountered); however, it is a bullet I am prepared to bite. I deny that to be not red is to be blue or green.²⁶ To those who find this example compelling, an explanation is owed: an account of why intuitions go awry in this particular case.

Note, to begin, that '*o* is not red' has one falsifier in common with '*o* is blue or *o* is green'. State *r* falsifies both expressions, since it verifies '*o* is red' and falsifies both '*o* is blue' and '*o* is green'.²⁷ This, arguably, is a metaphysically possible state of affairs; *o* is the type of object that could be red. The other falsifiers of '*o* is blue or *o* is green' are, equally arguably, metaphysically impossible. Object *o* cannot be both red and green, red and blue, or green and blue (at least not if this is to be understood as *o* being entirely red and entirely green at the same time). Impossible falsifiers are legitimate falsifiers, on the truth-maker approach, but they are often overlooked. When focusing our attention on the *possible* falsifiers, '*o* is not red' has the same falsifiers as '*o* is blue or *o* is green'. And it is for this reason, I maintain, that some may be inclined to accept 'To be not red is to be blue or green'. However, once our attention is broadened to the metaphysically impossible, the difference in falsifiers becomes apparent. It is not the case that to be not red is to be blue or green; what it is to be not red is to lack a particular phenomenal character—the phenomenology of redness.

If inference 9 is invalid, there ought to be cases in which it goes awry. If there are any such cases, they remain undiscovered. Every instance philosophers have discussed is one in which the inference from 'To be *F* is to be *G*' to 'To be $\neg F$ is to be $\neg G$ ' appears legitimate. Systems that deny the inference from 'To be *F* is to be *G*' to 'To be $\neg F$ is to be $\neg G$ ' fare worse than those that endorse it. Hence, that the present account licenses this inference is a mark in its favor.

A third difference concerns two axioms than Correia and Skiles license that I do not: 'To be *F* is to be $F \wedge F$ ' and 'To be *F* is to be $F \vee F$ ' (which are sometimes referred

²⁵ My thanks to an anonymous reviewer, who based this example on one in [23].

²⁶ This denial does not arise from the limited number of colors; I would deny that to be not red is to be blue or green or yellow, blue or green or yellow or orange, and so on.

²⁷ Here, as elsewhere, I assume that $r = r \sqcup r$ —that is, *r* is identical to the fusion of itself with itself.

to as ‘idempotence’). Correia and Skiles say little in their favor, aside from that they “[are] compulsory principles about generic identity” (pg. 11). In contrast, I deny both axioms. Recall that there is no guarantee that verification and falsification are closed under fusion; the fusion of two verifiers of ‘ Fa ’ may not verify ‘ Fa ’, and the fusion of two falsifiers of ‘ Fa ’ may not falsify ‘ Fa ’. In this case, a state that verifies ‘ $Fa \wedge Fa$ ’ would not verify ‘ Fa ’, while a state that falsifies ‘ $Fa \vee Fa$ ’ would not falsify ‘ Fa ’, and so both forms of idempotence fail.

There are deep, logical motivations for abandoning idempotence in this context—most notably discussed in [12]. Addressing these motivations in sufficient depth would take us far afield and require introducing further formalisms. Dorr argues that there is a reading of ‘To be grue is to be green and observed before time t or blue and not so observed’ that is true, yet ‘To be grue is to be green and observed before time t or bleen and not so observed’ is false. He accounts for this discrepancy by accepting a principle dubbed Only Logical Circles, according to which sentences of the form ‘To be F is to be F ’ are circular in an innocuous way, but nonlogical types of circularity are universally vicious. Once this principle is integrated with the λ -calculus (a standard way of representing logically complex predicates), it immediately follows that both forms of idempotence are false; it is straightforwardly possible to employ idempotence to derive nonlogically circular generalized identities. Those who endorse Only Logical Circles ought to abandon idempotence. Moreover, I note that the logic of generalized identity I advance is the most coarse-grained system consistent with Dorr’s principles.²⁸ And so, we arrive at the very same destination from distinct origins. While I am primarily motivated by the semantic considerations, Dorr is primarily motivated by logical ones. It is no accident, I think, that our systems so neatly align; I will say no more about these considerations here and direct those interested in the technical details to Dorr’s paper. I mention it only to note that the logic I advance is compatible with Dorr’s considerations, while Correia and Skiles’s is not.

There are plausible counterexamples to idempotence. Consider ‘To be $(F \vee G)$ is to be $(F \vee G)$ and $(F \vee G)$ ’. This is an instance of idempotence—in particular, of the conjunctive formulation—so if idempotence holds, sentences of this form ought to be true. One of the states that verifies ‘ $(Fa \vee Ga) \wedge (Fa \vee Ga)$ ’ is the fusion of a state that verifies ‘ Fa ’ with a state that verifies ‘ Ga ’, but there is no reason for this fusion to verify ‘ $Fa \vee Ga$ ’. For the sake of concreteness, consider ‘To be red or round is to be (red or round) and (red or round)’. One state of affairs that verifies that an object is ‘(red or round) and (red or round)’ is the state of it being both red and round (i.e., the fusion of a state that verifies that it is red with a state that verifies that it is round). This fusion need not verify that an object is red or round (at least on the exclusive formulation for disjunction), although it contains, as proper parts, two states of affairs that do verify that that object is red or round.²⁹

Much more could be said about the logic analysis, but it is my hope that the previous discussion suffices to motivate the present semantics over viable alternatives.

²⁸ I thank an anonymous reviewer for bringing this alignment to my attention—which Dorr demonstrates in the appendix.

²⁹ A parallel example could be constructed for the disjunctive form of idempotence. I thank an anonymous reviewer for this example.

§5. Conclusion. The focus of this paper has been a reflexive and symmetric notion of analysis, one that closely resembles an identity. However, I have briefly mentioned (and fully endorse) a reading of ‘To be F is to be G ’ that is irreflexive and asymmetric. Notably, it is possible to define this irreflexive and asymmetric notion (which, instead of a ‘generalized identity’, we might call a ‘definition’) in terms of the reflexive and symmetric notion on this account. Here, I distinguish this alternate reading with the use of the subscript ‘ df ’. On this reading, the ‘is’ of ‘To be F is $_{df}$ to be G ’ does not resemble the ‘is’ of identity; sentences of the form ‘To be F is $_{df}$ to be F ’ are universally false, and if ‘To be F is $_{df}$ to be G ’ is true, then ‘To be G is $_{df}$ to be F ’ is false. Before closing, allow me to briefly sketch an account of this reading.

The motivation for this approach is a thought that a definition of a phenomenon breaks it down into its constituent parts. While some may subscribe to this metaphorically, the current approach interprets it almost literally. Recall that verifiers and falsifiers are endowed with mereological structure: some are states with proper parts. The state of Ralph being a bachelor, for example, may be the composite of the state of Ralph being male with the state of Ralph being unmarried. And so, one can learn about the mereological structure of such a state by identifying its proper parts.

A natural thought is that a definition is a sentence of the form ‘To be F is $_{df}$ to be G ’ that satisfies two conditions:

1. ‘To be F is to be G ’ is a (true) generalized identity.
2. For all x , ‘ G ’ limns the mereological structure of the verifiers and falsifiers of ‘ Fx ’ more than ‘ F ’ does.

For example, ‘To be a mother is $_{df}$ to be a female parent’ is true just in case the states that make someone a mother are identical to the states that make her a female parent, and the term ‘female parent’ limns the structure of these states more than ‘mother’ does—presumably by demonstrating that they are fusions of the state of someone being female with the state of her being a parent. And ‘To be human is $_{df}$ to be a rational animal’ holds just in case the states that make something a human are identical to the states that make it a rational animal, and ‘rational animal’ reveals the structure of these states more than ‘human’ does—presumably by demonstrating that they are fusions of the state of them being rational with the state of them being an animal.

On this account, definition is irreflexive, asymmetric, and transitive.³⁰ The term ‘ F ’ never limns more structure than ‘ F ’ does, so sentences of the form ‘To be F is $_{df}$ to be F ’ are universally false. Similarly, if G limns more structure than F does, then F does not limn more structure than G does. So, if ‘To be F is $_{df}$ to be G ’ is true, then ‘To be G is $_{df}$ to be F ’ is false. And if ‘ G ’ limns more structure than ‘ F ’ does and ‘ H ’ limns more structure than ‘ G ’ does, then ‘ H ’ limns more structure than ‘ F ’ does. Because of this (and because generalized identities are transitive), if ‘To be F is $_{df}$ to be G ’ is true and ‘To be G is $_{df}$ to be H ’ is true, then ‘To be F is $_{df}$ to be H ’ is true.

This approach does not restrict the cognitive import of definition to linguistic facts. When one learns that ‘To be F is $_{df}$ to be G ’, she learns more than facts about the predicates ‘ F ’ and ‘ G ’; in particular, she learns about the mereological structures of

³⁰ The transitivity of the first condition follows from the transitivity of generalized identity. The transitivity of the second arises from the transitivity of ‘limns more structure than’. If G limns more structure than F , and H limns more structure than G , I assume that H limns more structure than F .

the states that make something *F* and *G*. In learning ‘To be a bachelor is_{df} to be an unmarried male’, for example, she thereby learns that the states that make it the case that someone is a bachelor are composed of the states that make that person unmarried and the states that make that person male. Definitions thus convey information about the world. A fully developed account of real definition ought to provide conditions under which a predicate limns more structure than another. I leave the development of these conditions for future work. I believe that this account of real definition is worthy of serious consideration, but will say no more about it here.

This theory of definition sheds light on the paradox of analysis, according to which it is doubtful that any analysis can be both informative and true. Substantive accounts are nontrivial: they convey information that agents may lack. But precisely because the content of an informative account differs from its object, that content cannot be the same as its object. The view that *being morally right* is_{df} *maximizing utility* is falsified neither by obligations to particular people nor by moral skepticism resulting from the plethora of unforeseen consequences, but rather by the mere fact that ‘maximizing utility’ conveys information that ‘being morally right’ does not. All quests for substantive analyses are futile, the paradox maintains, as their very substance undermines their truth.

Really, the paradox of analysis contains two puzzles. It is unclear, first, how *any* putative definition could be informative. But once that is resolved, it remains unclear why *all* putative definitions are informative—why it is that trivial definitions are universally excluded.

The present account suggests how both are to be resolved. Because definitions are a subset of identity claims, the first puzzle can be conceived as an instance of Frege’s puzzle about identity. The reason definitions can be informative is that identity claims can be informative. Just as ‘Hesperus is Phosphorus’ may convey substantive information, so too ‘To be a person is to be bound by the categorical imperative’ may convey substantive information. It is not committed to a particular *way* to resolve Frege’s puzzle—potential resolutions are numerous and varied. Rather, it suggests that any potential resolution could be applied to generalized identities in order to account for their substance. And the reason that *all* definitions are substantive is that they all reveal mereological structure. Trivial generalized identities are excluded precisely because—in virtue of their triviality—one term cannot limn more structure than the other. In order for ‘To be *F* is_{df} to be *G*’ to hold, *G* must reveal something that *F* does not—and it is for this reason that all definitions have substance.

There are advantages to this account of generalized identities that I failed to address. In particular, it offers insight into their modal status, the new riddle of induction, and the interpretation of physicalism.³¹ However, I do not rely upon arguments I lack the space to provide, but hope that the previous discussion motivates a conception of analysis in terms of truth-maker semantics and highlights some of the immediate logical implications of that position.

Acknowledgments. I would like to thank Fabrice Correia, Shamik Dasgupta, Michael Della Rocca, Catherine Elgin, Kit Fine, Jon Litland, Sun-Joo Shin, Seth Yalcin, and three anonymous reviewers for providing helpful comments on earlier

³¹ For discussions of these developments, see [14–16].

versions of this paper—as well as the attendees of the 2019 International Congress on Logic, Methodology and Philosophy of Science and Technology, Hamburg University's 2019 conference Truthmaker Semantics: Applications in Philosophy and Linguistics, the 2018 meeting of the Northern New England Philosophical Association, and the Canadian Metaphysics Collaborative's 2018 Conference on Essence and Language for their suggestions.

BIBLIOGRAPHY

- [1] Angell, R. (1977). Three systems of first degree entailment. *Journal of Symbolic Logic*, **42**(1), 147.
- [2] ———. (1989). Deducibility, entailment and analytic containment. In Norman, J. and Sylvan, R., editors. *Directions in Relevant Logic*. Dordrecht: Kluwer, pp. 119–143.
- [3] ———. (2002). *A-Logic*. Lanham, MD: University Press of America.
- [4] Bacon, A. (2019). Substitution structures. *Journal of Philosophical Logic*, **48**, 1017–1075.
- [5] Bacon, A. & J. Russell. (2019). The logic of opacity. *Philosophy and Phenomenological Research*, **99**(1), 81–114.
- [6] Barwise, J. & J. Perry. (1983). *Situations and Attitudes*. Cambridge, MA: MIT Press.
- [7] Caie, M., J. Goodman, & H. Lederman. (2019). Classical opacity. *Philosophy and Phenomenological Research*, **101**(3), 524–566.
- [8] Cameron, R. (2014). On the lack of direction in Rayo's the construction of logical space. *Inquiry*, **57**(4), 427–441.
- [9] Correia, F. (2016). On the logic of factual equivalence. *The Review of Symbolic Logic*, **9**, 103–122.
- [10] ———. (2017). Real definitions. *Philosophical Issues*, **27**(1), 52–73.
- [11] Correia, F. & A. Skiles. (2019). Grounding, essence and identity. *Philosophy and Phenomenological Research*, **3**, 642–670.
- [12] Dorr, C. (2016). To be F is to be G. *Philosophical Perspectives*, **30**(1), 39–134.
- [13] Dummett, M. (1991). *Frege: Philosophy of Mathematics*. Cambridge, MA: Harvard University Press.
- [14] Elgin, S. (2021a). Knowledge is closed under analytic content. *Synthese*, forthcoming.
- [15] ———. (2021b). The levels of the empirical sciences, forthcoming.
- [16] ———. (2021c). Physicalism and the identity of identity theories. *Erkenntnis*, forthcoming.
- [17] Fine, K. (1994). Essence and modality. *Philosophical Perspectives*, **8**, 1–16.
- [18] ———. (2006). Unrestricted quantification. In Rayo, A. and Uzquiano, G., editors. *Absolute Generality*. New York: Oxford University Press, pp. 20–44.
- [19] ———. (2013). A note on partial content. *Analysis*, **73**(3), 413–419.
- [20] ———. (2014). Truth-maker semantics for intuitionistic logic. *Journal of Philosophical Logic*, **43**(2–3), 549–577.
- [21] ———. (2016). Angelic content. *Journal of Philosophical Logic*, **45**(2), 199–226.
- [22] ———. (2017a). Survey of Truthmaker semantics. In Hale, B., Wright, C., & Miller, A., editors. *A Companion to the Philosophy of Language* (second edition). Chichester: Wiley, pp. 556–577.

[23] ———. (2017b). A theory of truthmaker content 1: Conjunction, disjunction and negation. *Journal of Philosophical Logic*, **46**(6), 625–674.

[24] ———. (2017c). A theory of truthmaker content 2: Subject-matter, common content, remainder and ground. *Journal of Philosophical Logic*, **46**(6), 675–702.

[25] Fine, K. & M. Jago. (2019). Logic for exact entailment. *The Review of Symbolic Logic*, **12**(3), 536–556.

[26] Frege, G. (1892). Sense and reference. *Zeitschrift für Philosophie und Philosophische Kritik*, **100**, 25–50.

[27] Fritz, P. (2019). Structure by proxy with an application to grounding. *Synthese*, forthcoming.

[28] Hwang, C.H. & L.K. Schubert. (1993). Episodic logic: A situational logic for natural language processing. *Situation Theory and Its Applications*, **3**, 303–338.

[29] Krämer, S. & S. Roski. (2015). A note on the logic of worldly ground. *Thought*, **4**(1), 59–68.

[30] Kratzer, A. (2007). Situations in natural language semantics. In Zalta, E.Z., editor. *Stanford Encyclopedia of Philosophy*. Stanford, CA: CSLI. Available at <https://plato.stanford.edu/entries/situations-semantics/> (accessed March 14, 2021).

[31] Lavine, S. (2006). Something about everything: Universal quantification in the universal sense of universal quantification. In Rayo, A. and Uzquiano, G., editors. *Absolute Generality*. Oxford: Oxford University Press, pp. 98–148.

[32] Linnebo, Ø. (2014). ‘Just is’-statements as generalized identities. *Inquiry*, **57**(4), 466–482.

[33] Parsons, C. (2006). The problem of absolute generality. In Rayo, A. and Uzquiano, G., editors. *Absolute Generality*. Oxford: Oxford University Press, pp. 203–219.

[34] Rayo, A. (2013). *The Construction of Logical Space*. Oxford: Oxford University Press.

[35] Rosen, G. (2015). Real definition. *Analytic Philosophy* **56**(3), 189–209.

[36] Tarski, A. (1933). The concept of truth in the languages of the deductive sciences. *Prace Towarzystwa Naukowego Warszawskiego, Wydział III Nauk Matematyczno-Fizycznych*, **34**(13), 172–198.

[37] Van Fraassen, B.C. (1969). Facts and tautological entailments. *The Journal of Philosophy*, **66**(15), 477–487.

UNIVERSITY OF CALIFORNIA, SAN DIEGO

9500 GILMAN DRIVE

LA JOLLA, CA 92103, USA

E-mail: elgin.samuel@gmail.com