Analyticity and Ontology

Louis deRosset

[forthcoming in *Oxford Studies in Metaphysics*]

April 1, 2013

Abstract

Analyticity theorists, as I will call them, endorse the *doctrine of analyticity in ontology*: if some truth ϕ analytically entails the existence of certain things, then a theory that contains P but does not claim that those things exist is no more ontologically parsimonious than a theory that also claims that they exist. Suppose, for instance, that the existence of a table in a certain location is analytically entailed by the existence and features of certain particles in that location. The doctrine implies that the table's existence requires nothing more of the world than that those particles exist and bear the features in question. Analyticity theorists have alleged that this idea may be used to defend controversial existence claims against a battery of objections. I argue that this style of defense fails, because the doctrine faces counter-examples. An existence claim may be analytically entailed by some truth and still report a substantial further fact.

Let's start with a story. A carpenter formulates a plan to make a table. She assembles some wood, shapes the pieces, and joins them according to the plan. What has the carpenter accomplished? How has she changed the way things are? She formulated and executed a plan, shaped some wood, and has more generally done what in the literature is called "arranging particles table-wise" in a certain spatiotemporal location L.[1] She has also produced a table in that location. We might summarize the import of this development thus:

EXISTENCE There is a table in L.

---

[1] The story I am telling makes some physical assumptions. In particular, it assumes that the wood from which the table is produced is made of particles, rather than, say, atomless gunk. I think the arguments of this paper would be unaffected if the physics of the story turns out to be false.

What's more, this table is, it seems, distinct from (and irreducible to) the microphysical effects of the carpenter's efforts. The table's distinctness from and irreducibility to those effects is attested by a difference in persistence conditions. The table could survive the passing of all of these effects together; under the right circumstances, the table would exist despite the fact that the particles which in fact make it up enter into a radically different arrangement. Thus, the carpenter has added to the furniture of the world. We might summarize the import of this claim by:

DISTINCTNESS The table in L is distinct from and irreducible to any congeries of particles.

At least, that's how matters appear to be.

Let's suppose we take these appearances at face value, and acknowledge that the table exists and is distinct from any congeries of particles that make it up. It would appear that we thereby add an additional entity to our ontology. The addition may be warranted by the arguments for **EXISTENCE** and **DISTINCTNESS**. If so, then we don't run afoul of Occam's Razor by "multiplying entities beyond necessity." We are, however, "multiplying entities." Other things being equal, the universe is less sparsely populated for containing both the table and its particles than it would be if it contained only the particles. At least, that's how it appears.

Authors whom I will call *analyticity theorists* argue that appearances are misleading because the existence of a table in location L is analytically entailed by the existence and arrangment of the particles there.[2] According to this view there is a specification of the existence and features of the particles which is such that it is analytic that if that specification is satisfied, then a table exists in the region where those particles are found. In short,

ANALYTICITY The existence of a table in L is analytically entailed by some appropriate specification of the arrangement of particles in L.[3]

---

[2]David Lewis has suggested that appearances are misleading for another reason: that mereological composition is ontologically innocent. Lewis's view is beyond the scope of this paper. See (Yi, 1999) and (deRosset, 2010, §6) for criticism of Lewis's contention.

[3]The qualifier "appropriate" is inserted to ensure that the specification in question does not *trivially* entail that the particles in L compose a table; for instance, 'the particles in L are arranged in such a way as to compose a table' would be an inappropriate specification. If pressed for a clearer account, we may assume that an appropriate specification is acceptable not only to the analyticity theorist, but also to theorists (*e.g.*, (van Inwagen, 1990)) who accept the existence of the particles but deny the existence of tables.

Something similar goes for other artifacts. In fact, claim analyticity theorists, the point generalizes to other entities, including works of art, rocks, mountains, and even properties, events, and numbers.[4]

Obviously the claim that the existence of a table in the relevant location is analytically entailed by the existence and arrangment of particles in that location relies on the notion of analyticity. Analyticity theorists are (among other things) thereby obliged to defend analyticity from the criticisms of Quine and others. They have taken this challenge up with enthusiasm. Their efforts won't be the focus of our discussion. Instead, let's give analyticity theorists the full measure of analyticity the view requires: assume that the Quinean criticisms can be put to rest, and that **ANALYTICITY** is true. We'll focus on the further claim made by analyticity theorists, that these analytic entailments buy an "ontological free lunch." On this view, accepting the existence of the table incurs no ontological commitments beyond those incurred by accepting that the particles exist and have the right arrangement. Admitting the existence of the table is "no addition to being." This idea is aptly expressed by Amie Thomasson:

> If claim ϕ analytically entails claim ψ, then competent speakers can infer the truth of ψ merely by knowing the truth of ϕ and knowing the relevant meanings of terms (and being competent reasoners). But if this is the case, then clearly ψ requires no more of the world for its truth than ϕ already required – sufficient truth-makers in the world for ϕ are also sufficient truth-makers in the world for ψ, they just make a new claim ψ true. (Thomasson, 2007, p. 16)

The claim, then, that will be our focus, the *doctrine of analyticity in ontology*, asserts a connection between analyticity and ontological parsimony:

DAO If P analytically entails the existence of certain things, then a theory that contains P but does not claim that those things exist is no more ontologically parsimonious than a theory that also claims that they exist.

This paper argues that **DAO** faces counter-examples. If the counter-examples stick, then analyticity does not buy an ontological free lunch. Even if the existence of a table is analytically entailed by the existence and arrangement of

---

[4]See (Thomasson, 2007), (Hale and Wright, 2001). Given **DISTINCTNESS**, the analyticity theorist requires the possibility of analytic entailment unaccompanied by reduction.

its particles, so long as the particles and the table are distinct, the table is "an addition to being." The objection is closely related to the familiar complaint that, according to **DAO**, we can, implausibly, "define things into existence."[5] Analyticity theorists have argued that this complaint is misplaced, denying that they are committed to the claim that, *e.g.*, the existence of any table *results from* the analyticity of any sentence. The analyticity theorist thinks, just as we do, that tables exist in virtue of the existence and features of their particles,[6] so no amount of linguistic legislation can conjure a table into existence. Even so, I will argue, there is a way of reconstructing this familiar complaint so that the underlying idea is correct. One source of the complaint is the idea that there is a sense in which the existence of a table in a given location is a *substantial further fact*, over and above the existence and arrangement of particles there and the rules governing 'table' talk; that's why, on this reconstruction of the complaint, the table can't be defined into existence. I will argue that this underlying idea is essentially correct. Despite **ANALYTICITY**, we have no reason to deny that the existence of a table in the relevant location is a substantial further fact in the relevant sense: **ANALYTICITY** does not imply that the table's existence is an "ontological free lunch."

I will begin by specifying in §1 some assumptions concerning analyticity and ontological parsimony that are required, both to properly understand the import of **DAO** and to carry out the arguments against it. Then in §2 I state the case against **DAO**. I discuss responses on behalf of the analyticity theorist in §§3-4. Finally, I close in §5 by sketching the picture of metaphysical inquiry suggested by the arguments.

1 Analyticity and Parsimony

Anayticity theorists claim that there is some specification of the existence and features of the particles in location L that analytically entails the existence of a table there. Let's suppose that $p_1, \ldots, p_n$ are the particles in L and abbreviate the specification in question by saying that $p_1, \ldots, p_n$ are "arranged table-wise." Then the analyticity theorist's claim is that

(1)   If $p_1, \ldots, p_n$ are arranged table-wise in L, then there is a table in L.

---

[5]See (Bennett, 2000, p. 56), for a contemporary attempt to press the objection.
[6]See (Thomasson, 2007, pp. 63-68) for a more extended response.

is analytic.

What does it mean to say that (1) is analytic? The notion of analyticity is notoriously difficult to specify clearly. For present purposes, we will understand the idea liberally: A sentence ϕ is *analytic* iff it is entailed by true sentences $\psi_1, \psi_2, ...$ such that failure to accept any ψ_n constitutes some measure of linguistic incompetence.[7] Call a sentence *foundationally analytic* iff it is true and failure to accept it constitutes some measure of linguistic incompetence. Given our understanding of analyticity,

(2)   If Mark Twain is a bachelor, then Mark Twain is unmarried

is analytic, since

(3)   No bachelor is married

is, plausibly, foundationally analytic in English. With the notion of analyticity in hand, we may define a notion of analytic entailment: a sentence ϕ *analytically entails* a sentence ψ iff the material conditional $(\phi \Rightarrow \psi)$ is analytic. Thus,

(4)   Mark Twain is a bachelor

analytically entails

(5)   Mark Twain is unmarried.

I hope that this specification is liberal enough that the analyticity theorist's claims concerning what is analytic in English are plausible. The specification lacks some of the alleged features of analyticity which theorists have found problematic. For instance, it does not require that analytic sentences are true *in virtue of* the rules governing our language; see (Boghossian, 1996). It also allows analytic truths that are surprising and informative. In fact, this specification allows that some analytic truths are are so long that no human being could

---

[7]We may not, however, infer from the fact that a speaker fails to accept one of the ψ's that she is linguistically incompetent *simpliciter*, nor that she fails to understand any expression of her language. The idea, instead, is that she falls short of utterly perfect linguistic competence. Presumably, all of us, even those of us competent in English, fall short of perfect competence. The other day, for instance, I was surprised to learn from the American Heritage Dictionary that 'meretricious' connotes something like 'ostentatious.' I hope nonetheless to be counted as competent *simpliciter* in English, and even competent *simpliciter* with 'meretricious'. What's more, nothing in the specification rules out the idea that these departures from perfect competence are theoretically motivated. Thus, the criticisms of (Williamson, 2007, ch. 4) seem to me not to apply.

entertain them. It does not require that analytic sentences follow from definitions of any sort, nor, more generally, from any specification of necessary and sufficient conditions for the extension of any term. It does not explicitly rely on any notion of synonymy, though undoubtedly some arguments to the effect that there is no such thing as synonymy may also undercut the presumption that there is such a thing as a foundationally analytic sentence – a sentence, that is, whose non-acceptance constitutes some measure of linguistic incompetence; see (Quine, 1951). It does not rely at all on any claim concerning the metaphysics of meanings, conceptual containment, or the like. It does not require or even suggest that analytic sentences are "about the meanings of our words" or "about language", rather than "about the world". It does not require that anyone who fails to accept an analytic sentence, or even a foundationally analytic sentence, fails to understand some of the expressions contained therein. More generally, it does not rely on any strong relation between accepting a sentence and understanding it; see (Williamson, 2007).

Our specification of analyticity is not, however, too liberal. No analytic sentence is false, since anything entailed by true sentences is itself true. And, it seems, there are plenty of true sentences that are not counted as analytic by the specification; for instance,

(6)   If Mark Twain is a bachelor, then Mark Twain is famous

is true but does not follow from truths whose acceptance is required for perfect competence in English. Likewise,

(7)   If Mark Twain is a bachelor, then the author of *Huckleberry Finn* is
        unmarried

is not analytic on this account, since failure to accept

(8)   Mark Twain = the author of *Huckleberry Finn*

does not constitute any measure of linguistic incompetence.

Our specification of analyticity is murky in crucial ways. First, it relies on a notion of entailment which has not been clearly specified. Second, it relies on the idea that failure to accept a certain sentence constitutes some measure of linguistic incompetence. I have emphasized (see n. 7) that being incompetent to some extent needn't imply being incompetent *simliciter*. Thus, the notion in question bears an indirect relation to everyday judgments concerning speakers'

linguistic incompetence and may for this reason be somewhat elusive.[8]   Third, something will need to be done about ambiguity, both lexical and structural, in order to deliver an answer to the question of whether (uses of) these sentences

(9)   If Twain drove a ship onto the ground at the side of a river, then he drove a ship onto a bank

(10) If Twain gave Crane a dollar, then he received money

are analytic. These are only examples of ways in which our conception of analyticity is unclear. These examples can undoubtedly be multiplied. I will be assuming, however, that the analyticity theorist can meet any reasonable demands for clarification that may arise. Nothing that follows will turn on any particular way of clearing things up.

Our specification of analyticity has an important epistemic upshot. In principle, speakers may come to know the truth of an analytic sentence by logical reasoning from sentences which, when considered, should be accepted by dint of linguistic competence. In this sense, competent speakers, on the basis of their competence and logical acumen, may reason their way to the conclusion that an analytic sentence is true.[9]   Thus, our specification of analyticity captures a methodological thread running through analyticity theorists' writing: analytic truths in my language may in principle be known solely by reflection on what my terms mean (together with an application of logical acumen).[10]

I think it's useful to have this rough and ready specification of a notion of analyticity in hand for the arguments that follow, but I don't think the arguments against **GAO** and **DAO** depend crucially on its correctness. In those arguments I will be working with cases in which analyticities are established by what I will call *linguistic stipulations*; linguistic stipulations are stipulations concerning what an introduced term is to mean. Every linguistic stipulation has a *content*, given by what is stipulated to be the case. For instance, if I make the linguistic stipulation

---

[8](Williamson, 2007, pp. 90-2,).

[9]In what follows, I will make no distinction between knowing the truth of an analytic sentence, and knowing the truth that sentence expresses; see (Donnellan, 1977), (Jeshion, 2001), and (Hawthorne and Manley, 2012) for discussion. In fact, I will indulge in use-mention sloppiness, since detailed attention to the distinction will not be relevant to the arguments and would be distracting.

[10]See especially (Boghossian, 1996) and (Thomasson, 2007).

(11) I hereby stipulate that: 'π' is a name that refers to the number (if there
is one) which, for any circle, expresses the ratio between the circle's
circumference and its diameter

then the content of my stipulation is given by the sentence that follows the colon.
A linguistic stipulation *succeeds* (alternatively, *is successful*) if, as a result of the
stipulation, its content is true.[11]   An assumption of these arguments is that,
under the right conditions, linguistic stipulations succeed, and thereby introduce
analyticities into the language. I will assume in particular that if a linguistic
stipulation succeeds, then the result of "disquoting" its content yields a non-
metalinguistic analyticity. Suppose, for instance, that I make the stipulation
(11). The content of that stipulation is a metalinguistic claim concerning the
name 'π'. My assumption is that, given the success of the stipulation, the
correlative non-metalinguistic claim

(12) If there is a number that expresses the ratio between the circumference
and diameter of a circle, then π is the number that expresses the ratio
between the circumference and diameter of a circle

is analytic. Even if the particular way I have chosen to articulate the notion
of analyticity is rejected, so long as sucessful linguistic stipulations give rise to
non-metalinguistic analyticities of the sort I have suggested, the arguments of
this paper will apply.

DAO implies that a theory is no less ontologically parsimonious for asserting
the existence of a table in a certain location than it is for claiming that the
particles in that location have the right arrangement. This claim crucially relies
on the idea of *ontological parsimony*. The expression 'parsimony' is clearly
metaphorical. What does this metaphor come to? That's a tough question, and
one I won't pretend to answer here in full.

We have two clues to guide us, however. First, there is the widely recognized
link between comparative parsimony and the application of Occam's Razor:
other things being equal, Occam's Razor favors a more parsimonious theory.
A vitalist theory that appeals to the existence of *elan vital* is less ontologically
parsimonious than a modern biochemical theory that does without it. Similarly,

---

[11]There are various ways in which a linguistic stipulation may fail to be successful. For
instance, the stipulation (11) fails if no term is added to the language, since in that case 'π'
would not be a name. That stipulation also fails, however, if 'π' is thereby introduced in the
language, but it isn't a name, or it is a name but refers to a boy.

Avogadro's hypothesis that oxygen molecules are composed of two oxygen atoms bonded together is favored by Occam's Razor over a competing hypothesis that they are composed of 34,000 oxygen atoms bonded together.[12] So, we have some insight into how Occam's Razor applies. The notion of parsimony on which these insights rely is the very notion employed by **DAO**.

Second, the analyticity theorist's claim about ontological parsimony is supposed to be an instance of the more general claim that if P is analytically entailed by some sentences, then P's truth "requires nothing more of the world" than does the truth of those sentences. So we can also use our insights into what the truth of a sentence or theory "requires of the world" to assess **DAO**. These insights may be hard to come by in some cases; the question of what Quantum Mechanics "requires of the world" is very difficult. But they are easy to come by in other cases, especially when what's at issue is the comparative question of whether one theory requires more of the world than another very similar theory, rather then the non-comparative question of what a theory requires of the world full stop. For instance, the truth of

(13) Snow is white

requires less of the world than the truth of

(14) Snow is white and grass is green.

I think it evident that (14) requires more of the world for its truth than does (13). But if an argument is needed there is one available. I will assume that requirements for the truth of these sentences are given by what their terms refer to and which properties their predicates express.[13] The truth of (14) requires that grass have a certain property, the property *being green*, which is expressed by the predicate 'green'. (13) does not impose this requirement. The requirements for (13)'s truth are satisfied or not according to whether snow has the right color; the requirements for (14)'s truth are satisfied or not according to whether snow has that color and also grass has a certain other color. For this reason, the requirements for (13)'s truth are satisfied and the requirements of (14)'s truth are not in worlds at which grass is purple but snow is still white.

---

[12]Avogadro's hypothesis was required to account for the ratios of volumes of oxygen gas and hydrogen gas to water vapor observed in combustion. See (Nolan, 1997) for discussion.

[13]Nominalists, who dispute the cogency of the idea that predicates express properties, are invited to substitute here their favorite replacement for the disputed idea.

So, the truth of (14) imposes a requirement on the world – that grass be green – that (13) does not.

Insights concerning the ontological parsimony of a given theory are just a special case: the parsimony of a theory is given by what the truth of the theory requires of the world *with respect to what there is*. So, a theory that asserts only

(15) There are more than 7 billion prime numbers

is more parsimonious than a theory that asserts

(16) There are more than 7 billion prime numbers and there are more than 7 billion human beings.

I think it evident that (16) requires more of the world than (15) with respect to what there is. But if an argument is needed there is one available. I assume that the requirements for the truth of these sentences are given by which properties their predicates express. The truth of (16) requires that there be more than seven billion instances of a certain property, the property *being human*, which is expressed by the predicate 'human being'. (15) does not impose this requirement. There are a wide variety of plausible views concerning what is required for the truth of (15), but no plausible view ties the satisfaction of these requirements to any particular count of the total human population. For this reason, the requirements for (15)'s truth are met and the requirements for (16)'s truth are not in worlds in which our arithmetical theories are true, but there are fewer than 7 billion human beings.

I will call the requirements of the world *with respect to what there is* a theory's *ontological commitments*.[14] Since ontological commitments are just a special case of requirements on the world,

GAO If P analytically entails Q, then $(P \wedge Q)$ requires nothing more of the world than does P.

---

[14]This use of the term "ontological commitment" differs from Quine's seminal treatment. According to Quine, the ontological commitments of a theory are given by what must be taken to be among the values of the theory's variables (under appropriate regimentation) for the theory to be true. Suppose, for instance, that we are given an appropriately regimented theory of artifacts whose truth requires that there be a table in L. According to Quine's treatment, the existence of such a table is among the ontological commitments of the theory. The present treatment, by contrast, is neutral. If the analyticity theorist is right, then the existence of the table needn't be among the ontological commitments of the relevant theory of artifacts.

is a generalization of **DAO** to which the analyticity theorist is also committed. In what follows, I will offer putative counter-examples to both **DAO** and its generalization **GAO**. The counter-examples will be of essentially the same sort. I hope thereby to show that the problems with **DAO** have nothing to do with any peculiarities of existential claims.

Here, in summary, is an upshot of **DAO**, as I am proposing to understand it. Suppose we are given a theory T. T has certain ontological commitments; that is, T imposes certain requirements on the world with respect to what there is. Let T^+ be the result of closing T under analytic entailment. According to **DAO**, T^+ and T have the same ontological commitments. Even supposing other things are equal, considerations of ontological parsimony do not favor T over T^+. The judicious use of Occam's Razor would never leave T and T^+ on different sides of the cut. For instance, T might be a theory exhausted by the claim we have abbreviated as '$p_1, \ldots, p_n$ are arranged table-wise in L' that details the existence and relevant features of particles in a particular location L. According to **ANALYTICITY**, (1) is analytic, so T analytically entails that there is a table in L. Thus, T^+ will contain the claim **EXISTENCE**, according to which there is a table in L. Nevertheless, T^+ has just the same ontological commitments as T; it requires nothing more of the world with respect to what there is than does T. More generally, T^+ requires nothing more of the world than does T.

In standard cases of analyticity, this upshot of **DAO** is very plausible. Suppose we are given a sociological theory according to which there are unmarried males. If we add to that theory the claim

(17) There are bachelors

we seem to have added nothing to the theory's commitments. This is a case in which the analytically entailed claim does not require the existence of anything distinct from the individuals already mentioned by the theory. But more interesting applications of **DAO** are also very plausible. Suppose we have a theory that says that Joe is a husband, but never explicitly mentions Joe's spouse. If we add to that theory the claim

(18) There is an individual married to Joe

then the resulting theory requires no more of the world than does the original theory, despite the fact that Joe's spouse was never explicitly mentioned by the

original theory. Thus, the existence of something may be required for the truth of a theory even if it is never explicitly mentioned by the theory. In both of these cases, the enriched theories simply slap labels onto things whose existence the more austere theories already required. The enriched theories may be more explicit about their commitments. They may usefully enable us to say more with fewer words. But they impose no further requirements on how the world is than their less explicit counterparts.

2   Three Stipulations

Despite the plausibility of **DAO** in standard cases of analyticity, I will argue that it faces counter-examples, in the form of analyticities introduced into the language as a result of successful linguistic stipulation. It is widely acknowledged that not every attempt at linguistic stipulation succeeds. Perhaps the most famous such example is 'tonk'.[15]   If the stipulation governing the use of 'tonk' were to succeed, then everything would analytically entail everything. Analyticity theorists have reacted to the possibility of bad stipulations of this sort by endorsing necessary conditions on successful stipulation that the stipulations governing 'tonk' and its ilk fail.[16]   I will argue that there are linguistic stipulations that plausibly pass all of the necessary conditions, and are intuitively successful to boot, but which still pose a problem for **DAO** and **GAO**.

There are a number of different conditions that theorists have proposed as necessary for successful stipulation. It will be useful to focus our discussion on just one, which is highly plausible and easily motivated. In general, it is plausible to think that a linguistic stipulation succeeds only if there is a consistent way to assign truth conditions to sentences containing the introduced term that makes the content of the stipulation true; call this *Stevenson's constraint* on linguistic stipulation.[17]   Stevenson's constraint is attractive. It is plausible to think

---

[15]A.N. Prior (1960) imagined that 'tonk' is stipulated to obey one of the introduction rules for disjunction and one of the elimination rules for conjunction:

$$P \vdash (P \text{ tonk } Q) \qquad\qquad (P \text{ tonk } Q) \vdash Q$$

[16]This strategy is followed in practice by analyticity theorists and articulated explicitly by (Hale and Wright, 2001, pp. 132-7). See also (Thomasson, 2007, pp. 171-2) and (Boghossian, 2003, p. 244).

[17]Stevenson (1961) originally pointed out that there is no truth table one may assign to 'tonk' that makes the proposed introduction and elimination rules valid.

that the function of linguistic stipulations is to introduce a new expression in such a way as to enable the interpretation of sentences that include it. If Stevenson's constraint is violated, then there is no coherent way to interpret sentences containing the erstwhile new vocabulary, since there is no coherent way to assign truth conditions to such sentences. So it seems that passing Stevenson's constraint is a necessary condition for a linguistic stipulation to fulfill its function.[18]

To illustrate the application of Stevenson's constraint, suppose I attempt to introduce the sentential operator 'it is verdantly the case that...', by attempting to stipulating that it be such that the following introduction and elimination rules are valid:

$$P \vdash \text{It is verdantly the case that } P$$

$$\text{It is verdantly the case that } P \vdash \text{Grass is green}$$

$$\text{It is verdantly the case that } P \vdash \Diamond P$$

This attempt at stipulation fails Stevenson's constraint. Consider the question of what truth value a sentence of the form 'it is verdantly the case that ϕ' should have when ϕ is true and grass is purple. The proposed elimination rule for 'verdantly' is invalid if the 'verdantly' sentence is true, and the proposed introduction rule is invalid otherwise. So, there's no way to coherently assign truth conditions to the operator that makes the content of the stipulation true.

A simple tweak to the stipulation, however, evades this response. Suppose 'it is verdantly* the case that...' is stipulated to be such that the following introduction and elimination rules are valid:

$$P \vdash \text{It is verdantly}^* \text{ the case that } P$$

$$\text{It is verdantly}^* \text{ the case that } P \vdash \text{Actually: grass is green}$$

$$\text{It is verdantly}^* \text{ the case that } P \vdash \Diamond P$$

---

[18]There are many other proposed constraints on successful stipulation. Belnap (1962) proposed that stipulations succeed only if the proposed rules of inference that are stipulated to govern the new vocabulary are *conservative* with respect to the old language, in the sense that the new rules do not permit the derivation of any sentences in the old language that were not already derivable. Hale and Wright (2001, pp. 132-7) endorse a weaker conservativity constraint and propose two additional constraints: (*i*) *Generality*: the stipulation should enable the interpretation of a wide enough range of relevant sentences; and (*ii*) *Harmony*: the introduction and elimination rules should not allow us to infer more (or problematically less) than our warrant for the premises allow us to infer.

The 'actually' operator here is the operator defined by Kaplan (1989): *actually:ϕ* is true at a world iff ϕ is true at the actual world. Notice that this stipulation appears not to fail Stevenson's constraint. Consider the question of what truth value a sentence of the form 'it is verdantly* the case that ϕ' should have when ϕ is true and grass is not green. It should have the truth value *true*, since, as a matter of fact, grass is green. So, we can coherently assign truth conditions to 'verdantly*' sentences.

If our stipulation concerning 'verdantly*' succeeds, then we face a counter-example to **GAO**. Suppose that we are given a theory T that includes the claim

(13) Snow is white

but does not mention grass at all. If our stipulation concerning 'verdantly*' succeeds, closing T under analytic entailment yields a theory T^+ that also includes the claim

(19) Snow is white and actually: grass is green.

The claims of T^+ would be true if T is; in fact, it is necessary that T^+ is true if T is. But that should not distract us from the important point here: T^+ requires more of the world than does T. We have already seen that

(14) Snow is white and grass is green.

requires more of the world than (13). In particular, (14), unlike T, requires for its truth that grass have the property *being green*. T^+ also imposes this requirement, so we have a counterexample to **GAO**. The argument here relies on the assumption that *actually:ϕ* requires no less of the world – the *actual* world, that is – than does ϕ. This assumption is powerfully plausible on its face, and is supported by the fact that the truth conditions for the Kaplanian actuality operator are specified by appeal to requirements on how the actual world is: what's required of the actual world for *actually: grass is green* to be true is just for the actual world to meet whatever requirements there are for *grass is green* to be true.[19] The 'verdantly*' stipulation, if successful, thus gives rise to a counterexample to **GAO**.

---

[19]Notice that *actually:ϕ* requires something different of the actual world for its truth from what it requires of a non-actual world. *Actually:ϕ* requires of the actual world that it meet whatever requirements there are for the truth of ϕ. The truth of *actually:ϕ* at a non-actual world w does not impose this requirement on w. The requirement on a non-actual world w is that it be such that some *other* world – the actual world – meet the requirement for the truth of *actually:ϕ*.

The analyticity theorist might reasonably complain that we have misapplied Stevenson's constraint. In assessing whether we can coherently assign truth conditions to 'verdantly*' sentences, we go wrong if we take for granted that, *e.g.*, grass is green, and go on to ask whether the sentence is true or false in some counterfactual situation in which grass has some color other than its actual color. Instead, we should consider what truth value to assign the sentence under the supposition that, *as a matter of fact*, grass is not green. In other words, when considering whether we can coherently ascribe truth conditions to sentences in the enriched language in the relevant situations we do what in the literature is called "considering the situation as actual."[20] Now, suppose that, *as a matter of fact*, grass is purple, and consider the question of what truth value should be assigned to

(20) It is verdantly* the case that: grass is either green or not green.

The proposed elimination rule for 'verdantly*' is invalid if the 'verdantly*' sentence is true, and the proposed introduction rule is invalid otherwise. So, there's no way to coherently assign truth conditions to the operator that makes the content of the stipulation true. Our stipulation fails this construal of Stevenson's constraint.

A slightly different stipulation evades this response. Suppose we stipulate that 'grassgreen' is to be a predicate that expresses the property *being green* if, as a matter of fact, grass is green, and *not being green* otherwise. Our stipulation will meet Stevenson's constraint, on the reading on which it requires that the inferences in question be truth preserving in all situations "considered as actual." Suppose that, *as a matter of fact*, grass is purple, and consider the question of what truth value should be assigned to

(21) grass is grassgreen.

Since grass is either green or not, the success of the stipulation for 'grassgreen' implies that this sentence be true.[21] This causes no problem for the interpretation of the sentence, however, since our supposition, together with the content of our stipulation, requires that 'grassgreen' express *not being green*. Similar

---

[20](Chalmers, 2002, p. 157).

[21]Technically, we may need to add the qualification "if grass exists" to (21) to get a sentence whose truth is implied by the success of the 'grassgreen' stipulation. Here and in what follows I will omit that qualification.

remarks apply to the supposition that grass is *as a matter of fact* green.[22]

So, Stevenson's constraint gives us no reason to think that the supposed stipulation concerning 'grassgreen' is unsuccessful. For this reason, it gives rise to apparent counterexamples to **GAO**. As before, let T be a theory of snow, and T^+ its closure under analytic consequence. T^+ would then contain

(22) Snow is white and grass is grassgreen.

T^+ is true if T is. In fact, we have, plausibly, an *a priori* guarantee that T^+ is true if T is.[23]   But that should not distract us from the important point here: T^+ requires more of the world than does T. Because grass is, as a matter of fact, green, the success of our 'grassgreen' stipulation implies that 'grassgreen' expresses the property *being green*. Once this fact is appreciated, I think it evident that T^+ requires more. But, if an argument is needed there is one available. I assume that the requirements for the truth of these sentences are given by what their terms refer to and which properties their predicates express. Because grass is green, the success of our 'grassgreen' stipulation implies that the predicate expresses the property *being green*. Thus, the truth of T^+ requires that grass have the property expressed by 'grassgreen', *i.e.*, the property *being green*. T does not impose this requirement. The requirements for T's truth are satisfied or not according whether or not snow has the right color and other relevant features. Thus, the requirements for T's truth are satisfied and the requirements of T^+'s truth are not in worlds at which grass is purple but snow is still white. So, the truth of T^+ imposes a requirement on the world – that

---

[22] Note that the stipulation governing 'grassgreen' is conservative, harmonious, and general. It's also, intuitively, successful.

[23] Boghossian (2003) considers the case of the stipulation that 'flurg' be governed by the following introduction and elimination rules:

$$x \text{ is an elliptical equation} \vdash x \text{ is flurg}$$

$$x \text{ is flurg} \vdash x \text{ can be correlated with a modular form}$$

Given that it is provable that every elliptical equation can be correlated with a modular form (this is the Taniyama-Shimura conjecture, proved by Andrew Wiles in the course of proving Fermat's last theorem), the analytic entailments supposedly introduced into the language by this linguistic stipulation are plausibly *a priori*. Boghossian argues (p. 244) that these supposed analyticities do not themselves give rise to an *a priori* entitlement to the Taniyama-Shimura conjecture, presumably because Wiles's sophisticated proof is not made accessible to us by the stipulation. Notice that this epistemological observation does not immediately bear on the plausibility of **GAO**, which is on its face independent of the accessibility to us of any instance of reasoning. Notice also that it is plausible to think that the stipulation governing 'grassgreen' plausibly does give rise to *a priori* entitlements to the alleged new analyticities, given the relative simplicity and transparency of the reasoning in question.

grass be green – that T does not.[24]

We have two apparent counterexamples to **GAO**, each appropriate to a different understanding of how to apply Stevenson's constraint. Neither of them concern existence claims, so they are not counter-examples to **DAO**. However, cases of essentially the same sort can be marshalled against **DAO**. Suppose we stipulate that 'priman being' is to be a predicate that expresses the property *being a human being* if, as a matter of fact, there are more than seven billion human beings, and *being a prime number* otherwise. The same sort of argument we gave in the case of 'grassgreen' applies here to show that the stipulation will satisfy Stevenson's constraint on the reading on which it requires that the inferences in question be truth preserving in all situations "considered as actual."[25] There is no problem consistently assigning an interpretation to 'priman being'.

Let T be a theory of arithmetic strong enough to prove that there are more than 7 billion prime numbers. If our 'priman being' stipulation is successful, then the result of "disquoting" the content of the stipulation

(23) If, as a matter of fact, there are more than 7 billion human beings, then
        something is a priman being iff it is a human being; otherwise,
        something is a priman being iff it is a prime number

is analytic. As a consequence

(24) If there are more than 7 billion prime numbers, then there are more than
        7 billion priman beings

is analytic.[26]   So, the closure of T under analytic consequence yields a theory

---

[24]If it is insisted that analytic truths be necessary (Kripke, 1980, p. 39), then we can tweak the example one last time, by mixing our last two cases. Suppose that 'it is verdantly** the case that...' is stipulated to have the following introduction and elimination rules:

$$P \vdash \text{It is "it is verdantly}^{**} \text{ the case that } P$$

$$\text{It is verdantly}^{**} \text{ the case that } P \vdash \text{Actually: grass is grassgreen}$$

$$\text{It is verdantly}^{**} \text{ the case that } P \vdash \Diamond P$$

The argument against **GAO** would then appeal to the claim, deployed in the discussion of 'grassgreen' above, that *actually:*ϕ requires no less of the (actual) world than does ϕ.

[25]The stipulation concerning 'priman being' also meets Belnap's proposed conservativity constraint, and satisfies Hale and Wright's Generality constraint. It satisfies Harmony, since, as I note below, there is plausibly an *a priori* guarantee that the theory T^+ (described below) is true if the arithmetical theory T (also described below) is true.

[26]Suppose there are more than 7 billion prime numbers. Either there are more than 7 billion human beings or there are not. If there are, then the priman beings are all and only the human beings and there are more than 7 billion of them. If there are not, then the priman beings are all and only the prime numbers and our supposition implies that there are more than 7 billion of them.

T^+ that contains

(25) There are more than 7 billion prime numbers and there are more than 7 billion priman beings.

As above, T^+ is true if T is, and there is plausibly an *a priori* guarantee that T^+ is true if T is. But that should not distract us from the important point here: T^+ clearly requires more of the world than does T. Since, as a matter of fact, there are more than seven billion human beings, the success of our 'priman being' stipulation implies that 'priman being' expresses the property *being human*, and is co-intensional with the predicate 'human', which also expresses that property. Once this fact is appreciated, I think it evident that T^+ requires more. But if an argument is needed there is one available. Since 'priman being' expresses *being a human being*, the truth of T^+ requires that there be seven billion instances of that property. T does not impose this requirement. There are a wide variety of plausible views on what the requirements for the truth of T come to, but no plausible view ties the satisfaction of these requirements to any particular count of the total human population. For this reason, the requirements for T's truth are satisfied and the requirements of T^+'s truth are not in worlds at which arithmetic is still true, but there are fewer than 7 billion humans. So, the truth of T^+ imposes a requirement on the world – that there be seven billion humans – that T does not.[27]

3 Quantifiers and the Counter-examples

So, we appear to have counter-examples to **DAO** and **GAO**, whether Stevenson's constraint is applied by "considering worlds as actual" or by evaluating the

---

[27]To get a case relevantly like the case of 'verdantly**', in which the analytic consequence in question is necessary, we might imagine stipulating that 'it is teemingly the case that ...' is an operator that obeys the following introduction and elimination rules:

$(P \wedge \text{there are more than 7 billion prime numbers}) \vdash \text{it is teemingly the case that } P$

It is teemingly the case that $P \vdash$ Actually: there are more than 7 billion priman beings

It is teemingly the case that $P \vdash \Diamond P$

Here, as with 'verdantly**', the argument appeals to the claim that a sentence of the form *actually:ϕ* requires no less of the actual world than does ϕ. More specifically, the claim required here is that *actually:ϕ* has at least as many ontological commitments as ϕ. It is not, I think, as clear that the 'teemingly' stipulation is successful as that the 'priman being' stipulation is successful; see §4.1 below for an argument that the original 'priman being' stipulation gives rise to a necessary analytic consequence that presents a counter-example to **DAO**.

truth of sentences at them in the more traditional way. How should analyticity theorists respond?

What they actually say seems not to help with the particular cases we have considered. Defenders of **DAO** have typically focused their efforts on the proper interpretation of quantifiers. Eli Hirsch (2002), for instance, defends **DAO** by appeal to a doctrine of *quantifier variance*, according to which there are multiple candidate meanings for the existential quantifier which are equally good for the purposes of ontology. On Hirsch's view, if the language is changed by the stipulative introduction of new vocabulary, so that new existential claims are apparently entailed, charity requires that our interpretation of the quantificational idioms shift to a new candidate meaning, so that the newly entailed quantificational sentences come out true. Imagine for the sake of illustration that ordinary English speakers originally speak a language in which they readily claim

(26) There is nothing other than particles in L.

Charity requires that we interpret them as deploying a candidate meaning for their quantifiers on which they speak truly. Imagine now that they stipulatively introduce the predicate 'table' into their language, so that they readily claim

(1)   If $p_1, \ldots, p_n$ are arranged table-wise in L, then there is a table in L.

On Hirsch's view, charity now requires that we shift our interpretion of the quantifiers: speakers are now to be interpreted as deploying a candidate meaning for their quantifiers on which this new sentence is true.

But this view seems not to help with the 'priman being' case. On the Hirsch response, that stipulation required that our interpretation of the existential quantifier shift so that, given the new quantifier meaning,

(27) There are seven billion priman beings

is entailed by our arithmetical theory. But if the relevant kind of entailment is *a priori* entailment, the old quantifier meaning will do just fine.[28]   No shift in our interpretation of quantifiers would help an analyticity theorist with the problem posed by the 'priman being' case. Despite the *a priori* entailment of T^+ by T, T^+ clearly has ontological commitments that T does not have.

---

[28]If instead the relevant kind of entailment is modal entailment, then we would have to appeal either to the example involving 'teemingly' (see n. 27), or to the argument of §4.1 below.

To take another example, Amie Thomasson (2007) has defended the existence of tables and other artifacts by arguing that existentially quantified sentences are uninterpretable without a contextually-supplied sortal, whose associated application conditions carry the relevant associated analytic entailments. Consider, for instance, the question of whether, *e.g.*, there is a table in location L. Thomasson argues that this question is uninterpretable unless some appropriate application conditions are associated with the sortal 'table'. Thomasson maintains that the question of whether there is a table in the relevant location is then answered by asking whether, in the relevant instance, those application conditions are satisfied. We may suppose that the application conditions for 'table' require that when particles of the right sorts exist and have the right features, then a table exists. Thomasson concludes that the proper semantic treatment of quantificational expressions in English thus implies that

(1)   If $p_1, \ldots, p_n$ are arranged table-wise in L, then there is a table in L

is analytic.

Thomasson's view does not provide a successful rebuttal of the apparent counter-examples. On Thomasson's theory, the interpretation of (27) requires that we specify the application and co-application conditions for 'priman being'. Nothing could be easier: because there are, as a matter of fact, more than 7 billion human beings, 'priman being' has the same application and co-application conditions as 'human being'. In fact, that's the reason why (27) incurs ontological commitments that T does not: (27) requires that there be more than seven billion loci in which the application conditions for 'human being' are satisfied. T, by contrast, doesn't impose any such requirement.

Thus, the particular semantic doctrines concerning existential quantification that are proposed by Hirsch and Thomasson do not appear to provide a way of avoiding the challenge posed by 'priman being'. Instead, their arguments, if successful, just show that (1) is analytic. The analyticity of (1) is something we have been supposing they are right about. The problem targeted by the 'priman being' and 'grassgreen' cases does not concern the question of whether (1) is analytic; the problem concerns the ontological significance of (1)'s alleged analyticity. So, Hirsch and Thomasson's arguments don't speak to the challenge posed by our apparent counter-examples.

There is, moreover, a more general reason for thinking that theses concerning the semantics of quantifiers won't provide an adequate defense of **DAO**.

There is a fundamental continuity between the case against **GAO** and the case against **DAO**. The continuity of our cases against **GAO** and **DAO** suggest that ultimately the defense of **DAO** must appeal to considerations that go beyond the proper semantic treatment of the existential quantifier. The proposed counterexamples to **GAO** do not involve the use of any existential quantifier, so **GAO** cannot be defended by appeal to any semantic thesis concerning only the existential quantifier. Since the proposed counterexample to **DAO** has essentially the same character, we should similarly expect a reply that goes beyond the semantics of quantifiers.

Suppose, finally, that some quantifier-focused defense succeeds for the 'priman being' case, but leaves the alleged counter-examples to **GAO** untouched. Then the plausibility of instances of **DAO** in focal cases of analyticity would no longer be enough to motivate **DAO**. The burden would be on the analyticity theorist to explain why we should expect the closure of a theory under analytic consequence to impose no new requirements on the world with respect to what there is, even though it may impose new requirements on the world in other respects.

4 Two Responses

What analyticity theorists have actually said does not appear to help with the apparent counter-examples to **GAO** and **DAO**. But the arguments that the stipulations in question present counter-examples to the analyticity theorist's view rely on two claims: (i) the stipulations generate certain analytic entailments; and (ii) the theory T^+ that results from closing a theory T under those analytic entailments requires more of the world than T. Thus, there are two potential avenues of response for the analyticity theorist. We'll discuss the propects for each avenue of response in turn.

4.1 Analyticity, Weak and Strong

The first potential response available to the analyticity theorist is to deny that the stipulations really give rise to the problematic analytic entailments that I have claimed. Let's focus on the 'priman being' stipulation. One way to pursue this strategy is to deny that this stipulation is successful. Making this claim presents the analyticity theorist with a challenge: articulate and motivate a

condition on the success of linguistic stipulations on which the 'priman being' stipulation is unsuccessful.

There is reason to think that this challenge will be difficult to meet. The 'priman being' stipulation does not just meet the constraints on successful stipulation that are found in the literature (see n. 18), it also seems to be utterly unexceptionable on its face. Thus, I suspect that the 'priman being' stipulation would present a *prima facie* counter-example to any otherwise well-motivated constraint on linguistic stipulation that would serve the analyticity theorist's purposes.

It would be unreasonable to rule out in advance the hypothesis of a lurking flaw in the stipulation; perhaps, after all, there's some subtle contradiction derivable from it. In the absence, however, of any indication of such a flaw, the smart money is on the success of the stipulation.[29]

There appears, however, to be a way of denying that the stipulation gives rise to the problematic analytic entailments while accepting the success of the 'priman being' stipulation: refine the explanation on offer of the notion of analyticity. The argument from the success of the stipulation to the falsity of **DAO** deploys a modal condition on the requirements for the truth of a theory: if it is possible for the requirements for T's truth to be met while the requirements for T^+'s truth are not, then T^+ requires more for its truth than T does.[30] The response is to suggest that the notion of analyticity in play likewise requires a modal specification: analyticities in the relevant sense must be restricted to

---

[29]The syntactic form of the 'priman being' stipulation is a conjunction of conditionals. I have found that this feature of the case gives some people pause. It is important to realize that this is an entirely typical syntactic form to use when making linguistic stipulations. For instance, if LeVerrier had been careful, he might easily have stipulated that 'Neptune' is to refer to the planet that causes perturbations in the orbit of Uranus if there is a unique such planet; and to refer to nothing otherwise. In fact, this syntactic form is so common in the specification of functions in mathematics that there are special notational conventions for abbreviating it.

[30]Notice that the constraint is stated in terms of the possibilities for meeting the requirements for the truth of the theories, rather than in terms of the theories themselves. This is because there are some theories that impose requirements for their truth on non-actual worlds different from the requirements they impose on the actual world. Thus, $(\phi \wedge$ *actually:*$\phi)$ and *actually:*ϕ clearly impose the same requirements on the actual world for their truth: to wit, what's required of the actual world for the truth of either is whatever is required of the actual world for the truth of ϕ. But, if ϕ is contingently true, then there is a possible world at which $(\phi \wedge$ *actually:*$\phi)$ is false and *actually:*ϕ is true. This is because the latter sentence does not impose the same requirements on the non-actual world at which ϕ is false as it does on the actual world; see n. 19. So, *actually:*ϕ is true at the non-ϕ world for a subtly different reason from the reason it is true at the actual world. Thanks to David Chalmers for pointing out this subtlety.

necessities. The analyticity supposedly introduced by the 'priman being' stipulation that causes all the trouble is the contingent truth:

(24) If there are more than 7 billion prime numbers, then there are more than 7 billion priman beings.

On this response, (24) is not analytic in the relevant sense, because it is not necessary. Thus, the stipulation is successful and gives rise to new entailments in the language, but it does not give rise to any untoward analyticities.[31]

Here is a way of carrying this idea out. Let's call the kind of analyticity we have been talking about up until now *weak analyticity*. Now define a stronger notion: ϕ is *strongly foundationally analytic* iff it is necessarily true and failure to accept it constitutes some measure of linguistic incompetence; ϕ is *strongly analytic* iff it is modally entailed by $\psi_1, \psi_2, ...$, such that every ψ_i is strongly foundationally analytic; and ϕ *strongly analytically entails* ψ iff the material conditional $(\phi \Rightarrow \psi)$ is strongly analytic. Intuitively, strong analyticity is just weak analyticity "plus necessity." These specifications ensure that only necessities are strongly analytic. Then the analyticity theorist can revise **DAO** to fit:

DAO$_\square$ If P strongly analytically entails the existence of certain things, then a theory that contains P but does not claim that those things exist is no more ontologically parsimonious than a theory that also claims that they exist.

I have argued that the 'priman being' stipulation gives rise to new weak analyticities in the language, including

(24) If there are more than 7 billion prime numbers, then there are more than 7 billion priman beings

As we have seen, (24) is no strong analyticity, since it is contingent. Thus, while (24) may present a counter-example to **DAO**, it presents no counter-example to **DAO**$_\square$.[32]

---

[31] Thanks to David Chalmers for suggesting the need to explore this avenue of defense in detail.

[32] This response comes with a cost: if **DAO**$_\square$ is to be useful for establishing that the existence of a table in location L is an "ontological free lunch" given the existence and arrangement of its particles, the analyticity theorist will need to establish that

(1) If $p_1, \ldots, p_n$ are arranged table-wise in L, then there is a table in L

is not only weakly analytic, but also necessary. See (Cameron, 2006) for discussion.

The analyticity theorist still faces a counter-example if three further premises are granted. The first is that, if ϕ is necessary, then

(28) $(\phi \Leftrightarrow \text{actually: } \phi)$

is strongly analytic. Instances of (28) in which ϕ is not necessary are weakly foundationally analytic, since failure to accept them indicates some measure of linguistic incompetence with respect to the Kaplanian actuality operator. They fail to be strongly analytic only because they fail to be necessarily true. If ϕ is necessary, then the instance meets this further condition for being strongly foundationally analytic. The second further premise is that if

(29) $(\phi \Rightarrow \psi)$

is weakly analytic, then

(30) $(\text{Actually:}\phi \Rightarrow \text{actually:}\psi)$

is strongly analytic. To see that this claim is true, suppose that (29) is weakly analytic, and so there are weakly foundationally analytic sentences $\psi_1, \psi_2, ...$ that jointly entail (29). But, for each of these weakly foundationally analytic ψ_i, *actually:*ψ_i is strongly foundationally analytic.[33] But then

$$actually{:}\psi_1, \qquad actually{:}\psi_2, \qquad ...$$

are strongly foundationally analytic sentences that modally entail (30).[34] The third premise we have already used in our discussion of 'verdantly*'; it is the

---

[33]*Argument*: (*i*) Since by hypothesis ψ_i is true, *actually:*ψ_i is necessarily true. (*ii*) Accepting ψ_i but failing to accept *actually:*ψ_i constitutes some measure of linguistic incompetence with 'actually'. Failing to accept ψ_i by hypothesis constitutes some measure of linguistic incompetence, perhaps with respect to other expressions in the language. Thus, failing to accept *actually:*ψ_i constitutes some measure of linguistic incompetence, with respect to either 'actually' or some other expression in the language. Argument (*ii*) may have to be qualified in order to take account of the fact that, *e.g.*, stringing a trillion 'actuality' operators onto the front of a weakly foundational analyticity may yield a sentence such that failure to accept it constitutes mortality rather than any measure of linguistic incompetence. Qualifications of this sort will not affect the argument in the main text, which involves very short sentences that may be taken in and assessed at a glance.

[34]*Argument*: $\psi_1, \psi_2, ...$ entail (29), so *actually:*ψ_1, *actually:*$\psi_2, ...$ entail

(31) Actually:$(\phi \Rightarrow \psi)$.

Because this sentence is necessarily true, the entailment is modally necessary. Further, it is strongly analytic that the relevant instance of the analogue for 'actually' of the axiom K for '$\Box$'

(32) $(\text{Actually:}(\phi \Rightarrow \psi) \Rightarrow (\text{actually:}\phi \Rightarrow \text{actually:}\psi))$.

is true.

plausible claim that the truth of *actually:ϕ* requires no less of the world than does the truth of *ϕ*.

With these three premises in hand, we can show that $\mathbf{DAO}_\square$ faces counter-examples of essentially the same sort as $\mathbf{DAO}$. Application of the first premise ensures that

(33) If there are more than 7 billion prime numbers, then actually: there are more than 7 billion prime numbers

is strongly analytic. Because

(34) If (there are more than 7 billion prime numbers ∨ there are more than 7 billion human beings), then there are more than 7 billion priman beings

is weakly (but not strongly) analytic, application of the second premise ensures that

(35) If actually:(there are more than 7 billion prime numbers ∨ there are more than 7 billion human beings), then actually: there are more than 7 billion priman beings

is strongly analytic. And, of course,

(36) If actually: there are more than 7 billion prime numbers, then actually: (there are more than 7 billion prime numbers ∨ there are more than 7 billion human beings)

is strongly analytic. Thus,

(37) If there are more than 7 billion prime numbers, then actually: there are more than 7 billion priman beings

is strongly analytic. Closing our arithmetic theory T under strong analytic consequence thus yields a theory T^+ which includes

(38) Actually: there are more than 7 billion priman beings.

Application of the third premise ensures that the original argument against $\mathbf{DAO}$ can be used to show that T^+ is less ontologically parsimonious than T. Thus, the success of our 'priman being' stipulation poses a problem for $\mathbf{DAO}_\square$, just as it does for $\mathbf{DAO}$.

4.2 Requirements and Truth Conditions

The second avenue of response accepts that our stipulations generate the relevant analytic entailments, but denies that those entailments are problematic. I have argued, for instance, that the theory T^+, which asserts

(22) Snow is white and grass is grassgreen

requires more of the world for its truth than the original theory of snow T from which it was obtained, on the grounds that, unlike the theory of snow, it requires of the world that grass have a certain color, green. But, a defender of **DAO** might urge, (25) does not impose any such requirement. I have claimed that the requirements for the truth of 'grassgreen' sentences like (22) are given in part by which property 'grassgreen' expresses. On the response we are now considering, requirements for (22)'s truth are not given in that straightforward and plausible way. That leaves us with the question of how they are given. So, adopting this response presents the analyticity theorist with a challenge: articulate and motivate an alternative view of the requirements for the truth of 'grassgreen' sentences on which (22)'s truth requires nothing more of the world than that snow is white.[35]

Such an alternative view is inspired by close attention to the content of the 'priman being' stipulation. According to the content of that stipulation, something is grassgreen iff it is green if grass is (*as a matter of fact*) green, and not green otherwise. This gives us a characterization of the requirements imposed by simple 'grassgreen' sentences like (21) that serves the analyticity theorist's needs: a sentence of the form 'α is grassgreen' requires for its truth that the referent of the term α be green if grass is, and not green otherwise. On this view, then, all that is required for the truth of

(22) Snow is white and grass is grassgreen

is that (*i*) snow be white and (*ii*) grass be green iff grass is green. It's plausible to think that this requirement does not go beyond the requirements for the truth of

(13) Snow is white

---

[35] The truth of (22) may also require the existence of grass; see n.21

since (*ii*) is trivially satisfied. If (*i*) and (*ii*) exhaust the requirements for the truth of (22), then the new analyticities introduced by our 'grassgreen' stipulation don't pose any problem for the analyticity theorist.[36]

The view of the requirements for the truth of 'grassgreen' sentences that backs this response is implausible. No doubt, *one* of the requirements for the truth (in the actual world) of sentences of the form 'α is grassgreen' is that the referent of α be green iff grass is. But there is *another* requirement for the truth of such sentences: the referent of α must have the property, *being green*, expressed by the predicate. If this is not also required for the truth of simple 'grassgreen' sentences, then it's a mystery why, *e.g.*,

(21) grass is grassgreen

is false at a world at which grass is purple, and so is green iff grass is. In general, our view of what is required of the world for simple 'grassgreen' sentences to be true ought to explain why (21) is true in some, but not all metaphysically possible circumstances. The view on which (21) requires only that grass be green iff grass is green doesn't pass muster. What's more, it's a mystery why in a world with green clover but purple grass

(39) clover is grassgreen

is true, even though clover is not green iff grass is green. The truth-conditions for simple 'grassgreen' sentences are thus left unexplained by the view of requirements for the truth of 'grassgreen' sentences on which the analyticity theorist's response draws.[37]

It would be unreasonable, however, to rule out in advance the possibility of offering a view of requirements that does better, while rendering the analytic entailments engendered by the 'grassgreen' and 'priman being' stipulations harmless to the analyticity theorist. Perhaps, after all, there's some subtle theory of the requirements imposed by such sentences that can do the relevant explanatory work. In the absence, however, of any such theory, the smart money is

---

[36]Thanks to Mark Moyer for pressing the need to explore this avenue of defense. Notice that this sort of response cannot be easily adapted to the case of 'verdantly*'. But I have already provided the analyticity theorist with a response for that case: the 'verdantly*' stipulation fails Stevenson's constraint on the appropriate way of applying it.

[37]It is plausible, perhaps, to think that the truth of (21) requires that grass be green iff grass is *actually* green; this is the view most naturally suggested by the content of the 'grassgreen' stipulation. But this requirement clearly goes beyond the requirement for the truth of the theory T of snow, as evidenced by the fact that the requirement is not satisfied in circumstances in which grass is purple but snow is still white.

on the claim that the truth of (21) requires, among other things, that grass be green.

5   Analyticity and Ontology

The stipulations we have imagined in this paper are silly. But, if the arguments concerning these stipulations are correct, the lesson they teach is serious. The fact that an existence claim is analytically entailed by some further claims about the existence and features of particles is no reason to think that the existence claim is "lightweight," and imposes no further cost with respect to ontological parsimony. The claim we imagined to have been stated in an enriched language by

(40)   There are more than 7 billion priman beings

is analytically entailed in that language by any suitably strong theory of arithmetic. It does not follow that the existence of more than 7 billion *homo sapiens* is "nothing over and above" the arithmetical facts, nor that our existence is an ontological "free lunch". Nor does it follow that whatever worldly condition or entity makes the theory true also makes it the case that there are more than 7 billion *homo sapiens*, nor that there is any truthmaker for the arithmetical theory that is also a truthmaker for the existence claim.

We should take no comfort in the thought that the stipulations we have discussed are contrived, while the analyticities concerning, *e.g.*, 'table' sentences occur naturally. The stipulations we have discussed are just a very explicit way of establishing conventional rules for using the terms in question. The fact that those rules are explicitly stated is what makes the examples seem contrived. We may not infer that the rules governing 'table' sentences, unlike the rules governing 'priman being' sentences, buy an ontological "free lunch" just because they arose in the give and take of natural language use, rather than being invented by a philosopher fishing for counter-examples. So, we should say the same thing about the existence of artifacts as we do about the existence of priman beings. We may not infer from the fact that **EXISTENCE** is analytically entailed by

(41) $p_1, \ldots, p_n$ are arranged table-wise in L

that its truth requires nothing more of the world than is required by (41). There may be other reasons for thinking so, but those reasons will need to appeal to

something other than analytic entailment.

Here, in broad brush strokes, is a way of thinking of the dialectic so far. We may think of analytic truths as having two features. First, they have what I will call *the trappings of analyticity*: the conventions of English guarantee that an analytic truth is entailed by certain sentences such that failure to accept any of these sentences constitutes some measure of linguistic incompetence. Second, they are true. The analyticity theorist contends that, if, *e.g.*,

(1)   If $p_1, \ldots, p_n$ are arranged table-wise in L, then there is a table in L

has the first feature, the trappings of analyticity, then its truth comes for free. I have argued, in effect, that this is wrong: granting that (1) has the trappings of analyticity, its truth may be a substantial further fact, imposing significant requirements on the world, including additional truthmakers for the consequent beyond those required for the truth of the antecedent.

Suppose that we were utterly convinced by these arguments against **DAO**. What view of the significance of analytic truths for metaphysical inquiry is thereby suggested? The arguments suggest that the utility of analyticity for investigation of extra-linguistic matters is virtually nil. If those arguments are correct, then

(24) If there are more than 7 billion prime numbers, then there are more than 7 billion priman beings

is analytic, and

(42) If there are more than 7 billion prime numbers, then there are more than 7 billion human beings

is synthetic. Further, the 'priman being' example can easily be generalized. Suppose $\Phi(F)$ is a claim that attributes a certain feature to the property expressed by the predicate F. (For instance, $\Phi(F)$ might be the claim that there are more than seven billion F's, or that Obama is F.) Let G be another predicate such that, as a matter of fact, $\Phi(G)$ is true. Then the stipulation

Let H be a predicate which expresses G if, as a matter of fact, $\Phi(G)$ and which expresses F otherwise

if successful, will give rise to the analyticity of

(43) If $\Phi(F)$, then $\Phi(H)$.

It's easy, then, to multiply examples so that, for every synthetic truth of the form *if $\Phi(F)$, then $\Phi(G)$* there is a corresponding analytic truth.[38]   It's hard to see how the distinction we thereby mark between these pairs of conditionals tells us anything interesting about the nature of extra-linguistic reality. It's tempting, then, to conclude that analyticity by itself has no significant metaphysical upshot at all, other than that analyticity requires truth. And, since the truth of an analytic sentence is not guaranteed by the fact that the sentence has the trappings of analyticity, giving in to this temptation means accepting that the trappings of analyticity by themselves have no significant metaphysical upshot.

   Suppose we give in to temptation. It's easy to come up with sentences that are true but lack the trappings of analyticity; any synthetic truth will fit the bill. But if the arguments of this paper are on track, then we should expect to find the converse situation, in which a sentence has the trappings of analyticity, but isn't true. We haven't yet come up with any sentences of this sort, but potential examples aren't hard to find. For the bulk of this paper, I have assumed **ANALYTICITY**, which says that the existence of a table in location L is analytically entailed by certain truths concerning the arrangement of the particles in L. Analytic entailment by truths requires truth, so this assumption implies that there is a table in L. That means that the assumption is dialectically inappropriate in the context of the debate between someone who denies and someone who affirms the existence of such a table. So, let's maintain the supposition that

(41) $p_1, \ldots, p_n$ are arranged table-wise in L

is true, but drop the supposition that any such table exists. Furthermore, we'll suppose that the participants in the debate over the existence of tables all speak ordinary English as it actually is, and that the analyticity theorist's claims concerning the norms governing the acceptance of 'table' sentences are correct. If no table exists in L, then

(1)   If $p_1, \ldots, p_n$ are arranged table-wise in L, then there is a table in L

---

[38]This recipe is just a generalization to predicates of Kripke's recipe (Kripke, 1980) for using linguistic stipulations governing names to generate contingent *a priorities*. The 'grassgreen' and 'priman being' cases are also inspired by those cases.

is not an analytic truth, since it's not a truth at all. Still the participants in the debate speak English, so (1) has the trappings of analyticity. Now, some of those sentences are false if no such table exists. So, if there is no table in L, we have sentences of exactly the sort we are looking for: accepting them is required for perfect linguistic competence, but they are false.

I think that many will find it dizzying to consider the possibility of a situation in which the semantic conventions of a language enjoin acceptance of a sentence which (in that very language) is false. There are three considerations that may mitigate the vertigo.

First, it is easy to imagine how someone might find herself in such a situation. Imagine that someone introduces the 'verdantly' operator into the language, and the operator is taken up by ordinary speakers of the language, who continue to use the operator and to apply its rules in accordance with the original stipulation. These language-users are particularly unreflective, however, and so never consider the possibility that, say, grass is not green, but purple. They teach their children to use the operator in accord with the inference rules, and correct them if they fail to accept

(44) If snow is white, then grass is green

much as a logic teacher might correct a student who fails to accept some slightly subtle logical truth, like

(45) if snow is white and grass is green, then grass is green and snow is white.

Thus, in this community, the conventions of the language, as embodied in both day to day use and instruction, appear to enjoin the acceptance of (44). Now just imagine that the members of this community wake up one day to discover that grass has turned conspicuously purple overnight. In this situation (44) has the trappings of analyticity other than truth (and, perhaps, necessity and *a priority*). The first person who thinks to herself,

(46) Snow is white, but grass is not green

is contravening the conventions of the language, as embodied in day to day use and instruction. Still, what she says is true. Given our assumption that (1) has the trappings of analyticity, the denier of the existence of a table in L says we are in fact in a similar situation. The conventions of our language, as embodied

in both day to day use and instruction, enjoin the acceptance of claims that turn out to be false.

Second, it is important to remember that suffering from some measure of linguistic incompetence is very widespread. In fact, it's almost certainly universal. By itself, failure to attain perfect competence in the use of one's terms does not constitute a very weighty charge against a theorist.[39]   Consider again the community of relatively unreflective speakers who use 'verdantly'. We have noted that on the morning that grass turns purple, the first person who thinks to herself,

(46) Snow is white, but grass is not green

is contravening the conventions of her language as embodied in day to day use and instruction. So what? Those conventions were established unreflectively in a way that had not taken account of the situation in which she finds herself. If this is the price of being right – of believing the truth – then it seems to me that price is right.

Third, it is important to realize that in such a case there may be and typically is some rearrangement of the conventions governing the language to accommodate the recalcitrant facts. There had better be: so long as the rules governing 'verdantly', 'grass', 'green', *etc.*, remain in force after the grass turns purple, speakers of the language are doomed to either error or some measure of linguistic incompetence. I have presumed throughout that new expressions may be introduced into a language in such a way that new conventions governing the use of those expressions may be established without altering the meanings of expressions in the old language. That means, for instance, that a new expression 'table' may be introduced into the language without altering the meanings of the expressions used to state the claim that $p_1, \ldots, p_n$ are arranged table-wise.[40] Similarly, I have assumed that the new operator 'verdantly' may be introduced into the language without altering the meanings of either

(13) Snow is white

or

---

[39]That is, it is not a very weighty charge so long as the measure of the theorist's incompetence is not so large that he fails to understand the terms he is using. See (Williamson, 2007, Ch. 4) for an argument that in the cases we are considering, the incompetence in question is not that large.

[40]Of course, the actual history of the expressions in English does not fit this description.

(14) Snow is white and grass is green.

In this sense, the introduction of new vocabulary extends the language rather than merely changing it.[41]   But, if adding "verdantly" to the language leaves the other vocabulary as is, then presumably removing it from the language does, too.[42]   Thus, when the speakers of the "verdantly" language wake up to purple grass, they aren't just stuck with a choice between error and some measure of linguistic incompetence. There's something they can do: pluck "verdantly", or at least the defective conventions governing its use, out of their language.[43]

Analyticity theorists face the challenge, posed by Quine, of accounting for the alleged fact that there are empirical data that might induce one to withhold assent to any sentence whatsoever, including erstwhile analyticities; the charge is, roughly, that each of our beliefs is in principle revisable in light of recalcitrant experience. In response, defenders of analyticity (Grice and Strawson, 1956) invoke the possibility of revising what one *means* (instead of revising what one believes) in the face of the new data. Thus, the meaning of the sentence rejected in one's new idiolect is not the same as the meaning of the sentence accepted in the old language. The present picture is neutral on whether the conventions in virtue of which sentences like

(44) If snow is white, then grass is green

have the trappings of analyticity should be taken to articulate the meaning of "verdantly." Other than that, there is just one thing that the present picture

---

[41]Thus, one way for the analyticity theorist to avoid the arguments above is to insist that the stipulations in question changed the meanings of other expressions in the language so that they no longer mean what they used to mean. The analyticity theorist might claim, for instance, that after the introduction of 'verdantly', the meaning of

(13)   Snow is white

somehow comes to express information about the color of grass; or he might claim that after the introduction of 'priman being', the arithmetical vocabulary of our theory of arithmetic somehow comes to express information about human beings. I take this sort of response to be *prima facie* implausible. What's more, it should provide scant comfort; see n. 42 below.

[42]Alternatively, if adding 'verdantly' to the language changes the meanings of the other vocabulary, so that, *e.g.*,

(13)   Snow is white

somehow comes to express information about the color of grass, then presumably removing 'verdantly' from the language can change those meanings back, so that (13) no longer conveys information about the greenness of grass, but is otherwise unaffected. The denier of the existence of tables should argue that, if this model applies to the case of 'table' talk, then English-speakers are obeying the conventions of the language when they affirm that there are particles arranged in the relevant way in location L, but what they say is (still) false.

[43]This is just one option. Speakers could, among other alternatives, decide to change what they mean by 'grass', by 'green', or even by 'is'.

adds: that sometimes revision of such conventions is not just convenient, but *required* by the twin demands to accept only truths and avoid linguistic incompetence. Carnap famously offered a list of pragmatic factors bearing on the question of which language to use:

> The purposes for which the language is intended to be used, for instance, the purpose of communicating factual knowledge, will determine which factors are relevant for the decision. The efficiency, fruitfulness, and simplicity of the use of [the language] may be among the decisive factors (1950, p. 208).

Our argument suggests that there is another decisive factor that is not merely pragmatic: whether the rules of use enjoin accepting falsehoods.

Quine and his followers are naturally interpreted as questioning the existence of a distinction between analytic and synthetic truths. If the arguments of this paper are correct, they should also have questioned its metaphysical significance.[44]

References

Nuel D. Belnap. Tonk, Plonk, and Plink. *Analysis*, 22(6):130–4, June 1962.

Karen Bennett. Composition, Colocation, and Metaontology. In David J. Chalmers, David Manley, and Ryan Wasserman, editors, *Metametaphysics: New Essays on the Foundations of Ontology*, pages 38–76. Oxford University Press, 2000.

Paul Boghossian. Analyticity Reconsidered. *Noûs*, 30(3):360–91, September 1996.

Paul Boghossian. Blind Reasoning. *Proceedings of the Aristotelian Society, Supplementary Volume*, 77:225–48, 2003.

Ross Cameron. The Contingency of Composition. *Philosophical Studies*, 136: 99–121, 2006.

Rudolf Carnap. Empiricism, Semantics, and Ontology. *Revue Internationale de Philosophie*, 4:20–40, 1950. Page references are to the reprint in (Carnap, 1956).

Rudolf Carnap. *Meaning and Necessity: A Study in Semantics and Modal Logic*. The University of Chicago Press, Chicago and London, second edition, 1956.

David J. Chalmers. Does conceivability entail possibility? In Tamar Szabó Gendler and John Hawthorne, editors, *Conceivability and Possibility*, pages 145–200. Oxford University Press, Oxford UK, 2002.

Louis deRosset. Getting Priority Straight. *Philosophical Studies*, 149(1):73–97, May 2010.

Keith Donnellan. The Contingent *A Priori* and Rigid Designators. *Midwest Studies in Philosophy*, 2(1):12–27, 1977.

H. P. Grice and P. F. Strawson. In Defense of a Dogma. *Philosophical Review*, 65(2):141–58, April 1956.

Bob Hale and Crispin Wright. *The Reason's Proper Study: Essays towards a Neo-Fregean Philosophy of Mathematics*. Clarendon Press, Oxford, 2001.

John Hawthorne and David Manley. *The Reference Book*. Oxford University Press, 2012.

Eli Hirsch. Quantifier Variance and Realism. *Philosophical Issues*, 12:51–73, 2002.

Robin Jeshion. Donnellan on Neptune. *Philosophy and Phenomenological Research*, 63(1):111–35, July 2001.

David Kaplan. Demonstratives. In Joseph Almog, John Perry, and Howard Wettstein, editors, *Themes From Kaplan*, pages 481–563. Oxford University Press, New York, 1989.

Saul Kripke. *Naming and Necessity*. Harvard University Press, Cambridge, 1980.

Daniel Nolan. Quantitative Parsimony. *The British Journal for the Philosophy of Science*, 48:329–43, 1997.

A.N. Prior. The Runabout Inference-Ticket. *Analysis*, 21(2):38–9, December 1960.

W. V. Quine. Two Dogmas of Empiricism. *Philosophical Review*, 60(1):20–43, January 1951.

J.T. Stevenson. Roundabout the Runabout Inference Ticket. *Analysis*, 21(6): 124–8, June 1961.

Amie L. Thomasson. *Ordinary Objects*. Oxford University Press, 2007.

Peter van Inwagen. *Material Beings*. Cornell University Press, Ithaca, NY, 1990.

Timothy Williamson. *The Philosophy of Philosophy*. Blackwell, 2007.

Byeong-Uk Yi. Is Mereology Ontologically Innocent? *Philosophical Studies*, 93: 141–60, 1999.