**What must exist for there to be laws of logic?**

The purpose of this piece is to spell out what I take to be the ontological commitments of logic.[1] Pretheoretically, we all take there to be truths that have a certain, distinct logical *flavour*. Examples include: "either it's raining or it's not raining", "all orange cats are cats", "nothing can be both living and not living at the same time". They are true, and the grounds for their truth apparently have something to do with logic. Now, the question is: what does this phenomenon of logical truth require of reality?

Many people take logic to have heavyweight metaphysical upshots. The archetype is the platonist. According to them, the logical truths reflect the laws that govern a realm of abstract entities—i.e. the propositions. These abstract propositions are supposed to exist necessarily, eternally, and independently of human affairs. They are not created by human language or cognition; on the contrary, the significance of our linguistic and cognitive activities depends on mirroring this abstract realm. The logical laws, then, are supposed to be facts about how these propositions pattern. For instance, it may be a law that no proposition and its negation can both be true (although, of course, some have disputed this). And according to the platonist, it is *because* this law governs the inhabitants of the abstract realm that my utterance of "all dogs bark and not all dogs bark" cannot express a truth. Some platonists might also add that these laws are normative; not only do they *describe* patterns of truth, but they also *prescribe* how we ought to reason.

If the platonist is the *most* ontologically committed in the metaphysics of logic, then what is the least? In the spirit of presenting caricatures, I'll pick on the *conventionalist* as the traditional opponent of platonism. (No doubt, there are numerous other positions. But I'm trying to present the archetypes that loom the largest in the popular imagination.) According to conventionalism, logical laws do not pick out features of a reality that existed prior to human activities. Rather, they are supposed to be reflections of the linguistic rules that were conventionally adopted in the development of human language. Languages are means of communication, and to communicate effectively, there must be certain ground rules that speakers mutually follow. On this line of thought, speaking a language is kind of like participating in a game, and games are defined, in part, by the rules that the participants must follow. If one is to play chess, then one can only move their bishops diagonally. If one moves the bishop in any other way, then one isn't really playing chess. Likewise, one might suppose that in order to count as participating in the English language game, one must treat sentences of the form *it is not the case that P* as derivable from sentences of the form *it is not the case that Q* and *either P or Q*. To fail to derive along these lines is to tantamount to opting out of speaking English. If so, then there's no deep metaphysical explanation as to why this rule is valid; it simply comes with the territory of speaking the English language. Finally, with this rule—and other such putative conventional rules—we can derive any sentence of the form *either P or not P*. So according to conventionalism, the law of the excluded middle—and the rest of the logical laws—are upshots of the conventional rules of language.

Platonism and conventionalism are two opposite extremes in the metaphysics of logic. My own view is that both of them are false. On the one hand, I believe that conventionalism fails to capture the way in which logical truth depends on worldly objects and their properties. However, the objects presupposed are just those that we already pre-theoretically take to inhabit the world. On the other hand, I do not believe that explaining logical truth requires us to posit an additional realm of

---

[1] Since it's a blog post, the usual provisions apply: I'm not attempting to do anything novel, I'm not trying to engage with the most recent literature, and so on. This piece is mostly for me. I'm the primary beneficiary because I can collect my thoughts.

language-independent *logical* entities. Logic is special in that the logical truths need not have a distinctly logical subject matter (contrary to platonism). And so the explanations of logical truth need not be traced back to a realm of abstract logical laws.

1. Conventionalism

I said that logical conventionalism looms large in the popular imagination, but nowadays this is mostly only true amongst non-professional philosophers. Conventionalism was once popular among empiricists, logical positivists, and Wittgensteinians, but ever since Quine's critique it has fallen out of fashion. I will not reiterate Quine's criticisms here because I do not favour them.[2] Instead, I want to focus initially on the difficulty in stating what exactly conventionalism amounts to.

The conventionalist claims that logical truths are made true by our linguistic conventions. They also claim that they reflect our decisions to use words according to certain rules. Sometimes they claim that the logical truths are "devoid of factual content" or that they are "not made true by the world".

So consider a putative logical truth: "Mars is either dry or not dry" (to borrow an example from Williamson and ignore non-classical objections for the moment). What would it mean to say that this is true by convention?

The first point to make is that there's a boring sense in which *all* sentential truths involve a conventional component: it is by convention that we use certain symbols in a given language to express their bestowed meaning. For example, it is a convention in English to use "Mars" as a name for the fourth planet from the sun. We could have used it as a name for something else, and other languages use different names. The same is obviously true for logical terms. We could have used a different word other than "and" to mean conjunction, and "and" could have been used to mean something else. So for logical conventionalism to be at all an interesting doctrine, it cannot simply be the point that our choice is conventional to use certain symbols to express certain logical meanings.

For an ordinary, non-logical truth, we can isolate the conventional aspect by conditionalizing its truth on the sentence's actual meaning. For example, we can say that: *given that "snow is white" actually means that snow is white (in English), it is true if and only if snow is white*. This shows us that there are two conditions that conspire to make "snow is white" true. The first is made true by our linguistic conventions: that "snow is white" actually means that snow is white (in English). The second is the language-independent, worldly fact that *snow is white*. The reason that "snow is white" is not conventionally true is that our linguistic conventions cannot affect the second, worldly factor.

We can similarly break up any logical truth into these two components. For instance, "Mars is

---

either dry or not dry" will be true given that (i) "Mars is either dry or not dry" means that *Mars is either dry or not dry* (in English) and (ii) Mars is either dry or not dry. Now, just like any logical truth, component (i) will be true due to the prevailing conventions of the English language. But that's the boring kind of conventionality that *all* sentential truths exhibit; it does nothing to support the conventionality of logic. If logical truths are to be conventional in any interesting sense, one must claim that the *second* factor—that *Mars is either dry or not dry*—is conventional. But is it?

Actually, it's very hard to make sense of the idea that the second factor is conventional because it doesn't ostensibly have anything to do with language. By all appearances, the second factor is about *Mars*, and it ascribes to Mars the property of being either dry or not dry. Because the second factor doesn't mention language (or propositions, or concepts, or truth), it's hard to see how it's any different from the non-logical case. That *Mars* is either dry or not dry appears to be just as much *about the world* (specifically, Mars) as the claim that *Mars is only ever dry.*

Some logical conventionalists—namely, A.J. Ayer—have claimed that a hallmark of conventionality is that a conventional truth is "devoid of factual content" and "says nothing about the empirical world". What I'm suggesting in the previous paragraph is that it is pretty patently false that "Mars is either dry or not dry" has nothing to do with the world. After all, it explicitly refers to Mars. As such, its truth depends, in some way, on how things stand with Mars. If Mars didn't exist, then "Mars is either dry or not dry" would suffer from reference failure. And given some standard assumptions about proper names, it would follow that "Mars is either dry or not dry" would fail to express a proposition or have a truth value.

This invites our first lesson about logical truth and ontological commitment. "Mars is either dry or not dry" is a logical truth (let's say, ignoring issues of vagueness), and it is (at least) ontologically committed to *Mars*. So at least one thing must exist for this logical truth to be true: namely, the planet Mars. Obviously, this lesson generalizes. "Venus is either dry or not dry" is a logical truth that's committed to Venus; "Jupiter is either dry or not dry" is a logical truth that's committed to Jupiter; and so on. We can also consider universally quantified logical truths "all things are such that either they are dry or not dry". What does this commit us to? Nothing in particular, but it is witnessed by all of the things that exist. The general lesson here is that logical truths *are* ontologically committing. But the ontological commitment found so far is just to the various things that we pre-theoretically knew to inhabit the world. A logical truth about Mars commits us to Mars, a logical truth about a people commits us to people, and a logical truth about tables and chairs commits us to tables and chairs. This much is true. But it does not *yet* demand any novel ontological commitment that stems specifically from the phenomenon of logical truth.

Now, returning to conventionality, if we grant that the fact that *Mars is either dry or not dry* (i.e. the second factor that explains the truth of "Mars is either dry or not dry") is foremostly about *Mars*, and not about language, then it's hard to see how it can be legislated true by our linguistic conventions. Presumably, we can legislate into existence conventional truths that are about human affairs; for example, I can legislate true that "Momo" will be a name, used by humans, as the name for my cat. But it's hard to see how I can legislate into existence facts about a planet that's millions of miles from Earth.

To add to this, consider the fact that *Mars is either dry or not dry* is a *disjunction*, and we cannot legislate-true either of its disjuncts. If Mars is in fact dry, then this would be so independently of our conventions. And if Mars is in fact not dry, then this would also hold independently of our conventions. Finally, a disjunction holds if at least one of its disjuncts holds. Therefore, it is a fact that *Mars is either dry or not dry* if either it is a fact that *Mars is dry* or is a fact that *Mars is not dry*, neither of which we have the power to make true by convention. This suggests that we do not have the power to make the

disjunction true by convention.[3]

Further evidence that *Mars is either dry or not dry* is not determined by conventions can be evinced by counterfactual reasoning. If something is made true by a convention, then it would not be the case if the convention did not prevail. For example, there is a convention in English of using the symbols "snow is white" to mean *snow is white.* But if that convention did not prevail—if "snow is white" had been used to mean something else—then it would not necessarily be true that *"snow is white" is true if, and only if, snow is white*. Well, consider these counterfactuals concerning the non-linguistic fact about Mars that it is either dry or not dry:

- If humans never had any linguistic conventions, then it would not be the case that Mars is either dry or not dry.

- If humans assigned different conventional rules to "or" and "not", then it would not be the case that Mars is either dry or not dry.

Both of these are false.[4] But counterfactuals like these are one of our best guides to cash out conventionality.

All of this goes to show that the second 'worldly' factor involved in explaining the truth of "Mars is either dry or not dry" is not a matter of convention. I've also noted that the first factor is a metalinguistic fact about meaning that is conventional in an uninteresting sense, no different from the way in which non-logical sentential truths rely on linguistic conventions. So we have so far failed to find any way in which logical truths are distinctly conventional.

The argument against logical conventionalism is not yet complete, however. There may be other ways to cash out conventionality. (It is hard to decisively refute a doctrine that is not clearly articulated.) I'll conclude by considering one other way of providing substance to the doctrine.

To this end, I'll invoke the foundational text in the project of analyzing conventionality: David Lewis's *Convention* (1969). In this book, Lewis provides an analysis of what it is for something to be a convention. The details of his account are too numerous to recapitulate here, but here's the gist of it: a convention (for Lewis) is a regularity in behaviour, sustained by mutual interests and expectations, *that has a possible alternative*. For example, there is a convention in North America of driving on the right side of the road. We all have an interest in avoiding head-on collisions, and so we regularly drive on the right side of the road and expect everyone else to do so. But we need not have chosen the *right* side to accomplish our goal. That choice is rather arbitrary. We could have just as easily fulfilled our mutual interests if we all drove on the left side.

Lewis's analysis captures the conventionality exhibited by the meaning facts of public languages like English. We all have a mutual interest in passing around information, such as the fact that snow is white. And we accomplish this if we all regularly utter the sounds "snow is white" when we wish to convey that snow is white. We could have easily accomplished this by some other means; instead of uttering "snow is white" we could have used a different string of sounds (as every non-English language

---

[3] These last two arguments come from Sider (2011) *Writing the Book of the World*, 6.5.

[4] I've heard that there is a philosophy professor who asks the following question to every interviewee for his school's graduate program: "If the word "white" were to mean green, then would snow still be white?". The answer, dear reader, is *yes*, because the consequent is about snow and whiteness, not about the word "white" (it is not metalinguistic). Moreover, the closest world in which the *word* "white" means something else is still a world in which snow is white.

does). But as long as we English speakers regularly use some common currency of sounds to convey the same information, we achieve the collective goal of communication.

Now, how can Lewis's analysis help us to explicate logical conventionalism? My idea is that the logical conventionalists may wish to claim that the meaning facts for logical truths are conventional *not only in the symbols we use to express them, but also in the meanings we wish to convey.* We've already observed that it's conventional (in the boring sense) that we use "and" to mean conjunction, rather than using some other symbol. But perhaps it's also conventional (in some profound sense) that we use "and" to mean *conjunction,* rather than some other equally valid, alternative logical operation.

To see what I mean, consider our systems of measurement. Surely it is conventional that we use *meters* as the unit for measuring distance (everywhere in the world except Liberia, Myanmar, and the United States). This can easily be explained using Lewis's analysis. We all have an interest in choosing *some* easy-to-eyeball distance for a unit for measuring mid-sized distances. But we didn't have to settle on meters. We could have used feet, or the length of my arm, or the height of the average human, or whatever. But we chose meters (rather arbitrarily). And once this choice was made, it's in all of our interest to stick to a common unit of measurement.

So let's say that "1 unit" is a name for the stand unit of measurement. "1 unit" thus displays *two* kinds of conventionality. First, it is conventional (in the boring sense) that the symbols "1 unit" were chosen to convey the information about the standard unit. And secondly, it's also conventional (in another, relatively more interesting sense) that "1 unit" happens to mean *one meter* (as opposed to any other alternative unit of measurement).

So now consider the logical case. I can imagine the logical conventionalist saying that there are several alternative schemes of inferential rules that we may associate with the logical terms. For instance, we happen to associate "and" with the following rules:

- (CI) A, B ⊢ A and B
- (CE1) A and B ⊢ A
- (CE2) A and B ⊢ B

And perhaps someone will say that this cluster of inferential rules gives "and" its meaning. Moreover, choosing *some* regular patterns of inferences involving "and" is to the mutual benefit of all English speakers. We all have an interest in "and" having a stable communicative effect, and so we expect others to continue inferring using "and" according to the given pattern.

But pursuant to conventionalism, the crux of this position is that this cluster of rules is just one, among many, equally-eligible clusters that could have been associated with "and". According to Lewis, a pattern is conventional only if there are alternatives. So if logic is conventional in the Lewisian sense, there must be a way in which our conventional choice of inferential rules is *arbitrary.* In that case, we *could* have chosen some alternative scheme of inferential rules that would have served our communicative purposes just as well. (This is just like how we could have chosen a different length to serve as the meaning of "one unit".)

According to the view that I'm describing, the truth of "Mars is either dry or not dry" is still divided up into two factors: a conventional linguistic component and a worldly component. We still explain its truth by saying that (i) "Mars is either dry or not dry" means that *Mars is either dry or not dry*, and that (ii) Mars is either dry or not dry. So we retain the idea that logical truths have worldly existential import that ultimately ground their truth. It's just that there are many alternative, equally-eligible schemes

of carving up the logical features of reality, and our chosen scheme is just one among many.

But now a problem arises. With logic (unlike units of measurement) there are strict independent standards for what makes a given rule *eligible* to govern a logical connective. According to one widely-shared conception of logic, logical rules must be *truth-preserving.*[5] If *<A and B ⊢ B>* is to constitute a governing rule for "and", it must be that if <A and B> is true, then <B> is true too.

This point about eligibility is made vivid by an example from Arthur Prior. Suppose that we introduced a 'logical' connective called "tonk" that is conventionally associated with the following two rules:

- (TI) A ⊢ A tonk B
- (TE) A and B ⊢ B

I suppose we *could* all agree to use this word and infer according to these rules. But if we did, we'd face an obvious disaster: the rules associated with "tonk" are not truth-preserving. From a true claim ("rabbits are mammals") we can use (TI) and (TE) to infer a false claim ("2 + 2 = 5"). In other words, the rules associated with "tonk" fail to meet the truth-preservation standard of logicality. So rules associated with "tonk" do not carve out a meaning that any genuine logical connective could have. Therefore, any sentence involving "tonk" cannot express a logical truth. Adopting the "tonk" rules as conventions would not be *logical*. The lesson, then, is that conventions, by themselves, do not create logical truth. In order for our logical vocabulary to capture logical truth, not just any meanings will do. The range of eligible meanings for our logical vocabulary is strictly limited.

This lesson from "tonk" is certainly awkward for the logical conventionalist, but it is not yet a knock-down objection. The conventionalist can agree that the range of admissible meanings for logical terms is limited, but insist that logical meanings are not uniquely determined without the help of conventions. Perhaps after the natural constraints (like truth preservation) narrow down the range of eligible systems of logical rules, there will still be some wiggle room. At this stage, I imagine that the conventionalist will bring up the various systems of alternative logic. Besides classical logic, there's also intuitionistic logic, relevance logic, paraconsistent logic, and so on. Perhaps the conventionalist will claim that each of these represents an eligible alternative meaning for the logical connectives, and it is just a matter of convention as to which one we actually adopt.

We have now entered into territory that is at the cutting edge of the philosophy of logic. The claim here is that there are multiple, equally viable systems of logic—a claim endorsed by logical pluralists—and that the adoption of one over any other is merely a matter of convention. In contrast, logical monists claim that there is only one true logic, and so there's only one cluster of rules that could capture the meaning of each logical term. It would be foolish of me to pretend to have a decisive answer to this debate. But for what it's worth, I favour monism. And if logical monism is correct, then this Lewsian version of logical conventionalism fails; the actual logical rules could not be chosen conventionally, for lack of viable alternatives.

My main reason for favouring monism is simply that the pluralist position doesn't appear to take the debates between the competing logics at face value. Ultimately, a system of logic is meant to capture a particular range of phenomena—of one thing following from another, and a thing being logically true.

---

[5] There are alternative logics that take other standards besides truth preservation as the independent eligibility criteria for rules of inference. But even if we consider those other standards, the point still stands that the resulting rules are not *arbitrary*, contrary to Lewisian conventionalism.

And the alternative logics make incompatible claims about what follows from what, and what is logically true.

For instance, classical logic claims that for any sentence P, ⌜P or not P⌝ is true, and intuitionistic logic denies this. To say that each respective position is right, but on different interpretations of the logical vocabulary, is to say that there's a sense in which ⌜P or not P⌝ is true for any P—it's just not one that's recognized by the intuitionist. But, in my view, to say this is actually quite uncharitable to the intuitionist's position. The intuitionist would not be happy to admit that the classical logician is right on *some* eligible meanings of "or" and "not" that they just so happen to have opted out of. Rather, they claim that there *isn't* any eligible meaning for "or" and "not" that would render ⌜P or not P⌝ a logical truth for every P.

A similar lesson can be gleaned from the debate between classical logicians and dialetheists. Classical logicians claim that there are no true contradictions (no true sentences of the form ⌜P and not P⌝), whereas dialetheists insist that there are. How should we reconcile the debate between them? The pluralist claims that the classical logician and the dialetheists are both right according to different, but equally legitimate interpretations of "and" and "not" (and perhaps "true"?). It's just that the classical logician uses "and" and "not" differently than the dialetheist. Now I hope the reader will agree that this *isn't* really a charitable way of reading the classical logician's position. The classical logician really wants to insist that there are *no* legitimate interpretations of "and" and "not" that would permit ⌜P and not P⌝ to come out as true.

The point is: the pluralist's attempt to make each side right actually ends up doing an injustice to one side. Saying that both sides are right in their own idiolects is often tantamount to saying that (at least) one side is simply mistaken.

In fact, I think that this lesson generalizes. There is a certain kind of philosopher (and innumerable philosophy students) who are prone to saying that debates in metaphysics are merely semantic or verbal. (Indeed, these are the very same people who are likely to be attracted to logical conventionalism.) They see metaphysicians making opposing claims and they figure that the debate can be resolved by interpreting both of their claims charitably, so that one metaphysician is speaking truly in one idiolect and the other metaphysician is speaking truly in another. Once they say this much, they might as well say that the dispute can be settled one way or another if we all *conventionally* adopt a single, stable idiolect. The problem with this cheap way of resolving metaphysical debates is that it pretty much always fails to do justice to the initial debate. To say that both metaphysicians are right in their own way is to beg the question against (at least) one of them.

So what does all this mean for logical conventionalism? If I'm right, then it's hard to say that there are multiple, equally legitimate contenders for the meanings of our logical terms. (More likely, there is just one set of meanings: i.e. the ones that carve out the truth-preserving forms of inference.) And if it's hard to say this, then it's equally hard to say that the choice of meanings for our logical terms is conventional in the Lewisian sense. Lewisian conventionality requires that the choice of meanings for our logical terms is *arbitrary* (in a certain sense). But the meanings for our logical terms are patently *not* arbitrary; they are constrained to underwrite the facts of what follows from what.

## 2. The subject matter of logic

Before we can dive into the ontology of logic, we need to attend to two preliminary matters. First, we need to get clear on what makes a claim, sentence, theory, etc. ontologically committing. Under what

circumstances does accepting the truth of one thing commit us to the existence of another? And secondly, we need to lay down the subject matter of logic, so that we have some materials to start with.

The problem is that both of these matters are far beyond any paper. The first is the central question of metaontology and the second gets to the heart of the philosophy of logic. Indeed, it's very difficult to specify the subject matter of logic while remaining theory-neutral. Different philosophies of logic make different claims on what logic is all about.

As for metaontology, I'll shortcut the discussion by stating, somewhat dogmatically, what I believe. Like Quine, I take a theory's ontological commitments to be revealed by its use of existential quantifiers. Hence, any theory containing statements of the form "there is something that is F" is thereby committed to the existence of Fs. But unlike Quine, I also take a theory's ontological commitments to be revealed through its use of proper names.[6] Hence I claim that "Mars is either dry or not dry" commits us to the existence of Mars.

These two positive claims about commitment might sound truistic, but the significance, for us, is in their austerity. I don't think that there's any other way for a theory to be ontologically committing; ontology is revealed solely through objectual quantifiers and proper names. If this is right, then there is no easy route to establishing the existence of the exotic entities of platonic metaphysics. We can't argue, for example, that universals must exist because Mars is red and redness is a universal. In the sentence "Mars is red", "red" does not occupy an existentially-committing position. There may be other, more roundabout ways of arguing for the existence of universals. But this flat-footed way is not one of them.

Now, as for logic, the most theory-neutral way to begin to identify the subject matter is by listing off a number of putative logical truths and logical entailments. Uncontroversially, "it's raining" follows from "if there are clouds in the sky, then it's raining" and "there are clouds in the sky"; and uncontroversially, "all orange cats are cats" is logically true. The trouble begins when we try to say what all of the uncontroversial entailments and logical truths have in common.

There are roughly three approaches to addressing this problem. There's the Tarskian (semantic, model-theoretic) approach, the modal (platonist) approach, and the syntactic (proof-theoretic) approach. Each of these is naturally paired with some fairly big-picture assumptions about the nature of meaning and the function of language, so arguing over them will take us straight to the fundamental questions of the philosophy of language. I'll take them in reverse order.

According to the syntactic, proof-theoretic approach, logic is fundamentally about certain rules of inference that govern the use of the logical vocabulary. (We've seen something of this idea already in the attempt to flesh out a Lewis-style version of logical conventionalism.) The idea is that the very meaning of the logical terms, like "and", cannot be explained any more fundamentally than citing the inference patterns in which they figure (i.e. conjunction-introduction, conjunction-elimination, etc.). These inference patterns are to be specified syntactically (i.e. in terms of the *form* of the premises and conclusion), and not in terms of some more fundamental semantic constraint, like truth preservation. Perhaps these rules are constitutive of the human "language game", or they are grounded in the nature of reason, or whatever.[7] This idea pairs well with the idea that meanings are fundamentally cashed out in terms of patterns of use (the so-called "use theory of meaning"). And according to this picture, language

---

[6] Quine has his tricks for simulating this, but we need not delve into it.

[7] One strength of this approach is that it has a ready-made explanation of the alleged *normativity* of logic. Namely, that logic is normative because it is a matter of following the rules that are constitutive of reason or the human language game.

I'll say this later on, but I don't myself think that logic is normative in any interesting sense. So this apparent advantage of the syntactic approach does not move me much.

is to be understood on the model of a game that has certain rules about what to say (or think) when and where; it has less to do with the expression of thoughts or representing the facts of the world.

This approach to the subject matter of logic is ontologically committed only to whatever explains the existence of these rules. If the rules arise from human conventions, then it's thoroughly naturalistic. If the rules are constitutive of rationality itself—they are the "laws of thought" (whatever that means)—then the view is committed to rationality having an inherent structure. Or if the rules are laid down by God—a divine convention-arbiter—then this view will be theistic.

I've already given one reason to disfavour this view in my discussion of Lewisian conventionalism. Namely, that syntactically-defined rules of inference should only govern the use of a logical expression *if* they are truth-preserving. This suggests that *truth-preservation*—a semantic notion—is more fundamental to the meanings of logical terms than syntactically-characterized rules of inference. Besides this, I would also add, on a personal note, that I just don't really buy the idea that language is like a game. On the contrary, I believe that a central function of language is *representing things in the world*. Communication is not just a matter of patterning our speech to make appropriate sounds back and forth with our interlocutors, like what chat-GTP is trained for. Rather, the point is to say *true things* about the objects in our environment using symbols that stand for them. If this is the right picture of language, then the role of logic is to facilitate this function by describing what must be true given that other things are true. But this role is inherently semantic. It cannot be captured by a purely syntactic approach.

The next approach to logical consequence and logical truth appeals to the notion of necessity. The idea is that logical consequence amounts to necessary consequence, and perhaps logical truth can be cashed out as necessary truth. This idea might be paired with the platonic idea that the arena of logical truth consists of necessarily existing abstract objects, but we will keep these ideas separate. The main idea that we'll consider is that necessity is a mark of logicality. So, for instance, what makes "it's raining" a *logical* consequence of "if there are clouds in the sky, then it's raining" and "there are clouds in the sky" is the fact that it is *necessary* that the former is true if the latter is true.

Despite its esteemed tradition in the early 20th century, I don't think anyone nowadays would think that necessity is sufficient for logicality. That there are necessary truths that are not logical truths is one way of putting Saul Kripke's point about the necessary a posteriori (assuming that logical truths are a priori). One famous example is water being necessarily $H_2O$. That water is $H_2O$ is arguably necessary since it's in the essence of water to be $H_2O$, but this fact is patently not a *logical* truth. That water is necessarily $H_2O$ follows, in part, from empirical information about the actual chemical composition of water.

So necessity is not sufficient for logical truth. It is also not sufficient for logical consequence. It is necessary that Hesperus is a planet if Phosphorus is a planet (because Hesperus and Phosphorus are one and the same planet). But "Hesperus is a planet" is not a logical consequence of "Phosphorus is a planet"; it follows from the extra-logical fact that Hesperus is Phosphorus.

While few people nowadays think that all necessary truths are logical, there are still a number of people who think that all logical truths are necessary. Indeed, this is even codified in the systems of normal modal logic through the rule of necessitation. Basically, the rule of necessitation says that if a given formula is a theorem of propositional logic, then its necessitation is a theorem of (normal) modal logic. For instance since "P $\lor$ ~P" is a theorem of PL, "□(P $\lor$ ~P)" is a theorem of any normal system of modal logic.

Now, it makes sense to include this rule in a formal system where the purpose is to model the logical inferences that involve specific types of logical operators (i.e. "necessarily" and "possibly"). But

we shouldn't take this to suggest that it's a *general* requirement on logical truth. After all, there are other kinds of logical truth besides those that are captured by propositional and modal logic. And arguably, there are instances of truths that are *logically* true without being *necessarily* true.

To see this, consider the logic of "actuality". Arguably, "actually" is a logical operator. After all, it's just as amenable to formal treatment as "necessarily" and "possibly", and the study of the logic of actuality is just as fruitful as ordinary modal logic. Moreover, there are theorems of the logic of actuality whose necessitations are false. Case in point: "It is raining if and only if it is actually raining". Think of it this way. Choose a certain world to designate as the actual world; it doesn't matter which one. From the perspective of that world, "It is raining if and only if it is actually raining" will be true no matter what. So no matter what the actual circumstances, this sentence is true; hence, it's logically true. But its necessitation will be false given that "it's raining" is contingent. Because whether or not it's raining in the actual world, there will be some other world with the opposite weather; hence, "It's raining if and only if it's actually raining" will be false with respect to that world.

Anyway, the upshot is that *necessity* is neither necessary nor sufficient for logicality. Whatever logic is, it cannot be defined in modal terms.

This brings us now to the final view of logical truth and consequence: the Tarskian view, which I favour. Consider again the sentence "Mars is either dry or not dry". What makes this a logical truth (again, assuming that it is, ignoring vagueness)? According to the Tarskian, this qualifies as a logical truth because the contributions of the logical terms "or" and "not", along with the sentence's syntactic structure, suffice to guarantee its truth. We can put this point by saying that it doesn't really matter what the other, non-logical parts of the sentence mean. The name "Mars" could, for all we care, refer to the big toe on my left foot, and the predicate "dry" could express the property of being a prime number. No matter how we interpret the non-logical parts, the sentence will always come out as true. In this sense, we can say that the sentence is true by virtue of the meanings of the logical terms and its form; or better: it's true by virtue of logic.

The Tarskian account of logical consequences is a continuation of the same thought. Accordingly, S is a logical consequence of a set of sentences E if, and only if, no matter how the non-logical terms of S and E are interpreted, S is true whenever every sentence of E is true.

Both of these accounts explicate the idea that logical truth and consequence can be traced to the meanings of the logical vocabulary and syntactic structure. Evidently, a lot of heavy lifting is done by dividing terms up into the "logical" and the "non-logical". So advocates of this view owe us a story as to what makes a term a *logical* term. But this challenge can be met. (I've got this idea from Gila Sher's *Epistemic Friction* (2016).) Basically, for the Tarskian, a term is a logical term if its meaning is one that's indifferent to what individuals there are, and what their distinctive features are.[8] There just have to *be* objects and properties, and then no matter what they are, it will be meaningful to speak of conjunctions, disjunctions, quantifications, and so on. This contrasts with non-logical terms, whose meanings depend on the distinctive features of objects. "Mars" is a non-logical term as it refers to a particular planet; it would have no referent if this planet did not exist. "Dry" is a non-logical term, because its application to objects depends on distinctive features of them—i.e. their amount of moisture.

Putting these two ideas together, we get the following accounts:

---

[8] Kant writes that "general … logic… treats of understanding without any regard to the difference in the objects to which the understanding may be directed (Kant 1781/7: A52/B76). Frege writes that "pure logic… disregard[s] the particular characteristics of objects" (Frege 1879: 5).

- S is a logical truth if, and only if, it is true under any interpretation of the terms of S whose actual meanings depend on the distinctive features or identities of objects.

- S is a logical consequence of E if, and only if, under any interpretation of the terms of S and E whose actual meanings depend on the distinctive features or identities of objects, S is true whenever every sentence of E is true.

You might have noticed that there's a tension between this account of logical truth and my discussion in the previous section. When I argued against conventionalism, I argued that logical truths have a "worldly factor" and that they are not really "about language". But the Tarskian account of logical truth is overtly metalinguistic. So one might wonder: does the Tarskian account ignore the worldly factor of logical truth?

I would argue that it doesn't. To explain why any given logical truth is *true*, we appeal to its actual truth conditions, which involve a worldly factor. For instance, explaining why "Mars is either dry or not dry" is true involves citing the fact about *Mars* that it is either dry or not dry. And then to explain why this truth has another feature, *logicality*, we give the Tarskian account which is more metalinguistic. In that case, we cite the additional fact that the sentence would *remain true* under any alternative, non-actual interpretation of "Mars" and "dry". We'd say that "Mars" could refer to any object, and "dry" could apply to any range of objects, and the sentence will be true no matter what. This metalinguistic story serves to explain its status as a *logical* truth, which is distinct from explaining why it's *true* (simpliciter). It's also not purely metalinguistic. The Tarskian account invokes interpretations which are essentially functions from names to objects and predicates to sets of objects. So it makes at least one metaphysical demand of the world: that the world be composed of *objects* (as opposed to, say, events, facts, undifferentiated bundles of tropes, or whatever). That the world is divided into objects is an assumption about its metaphysical structure, and it is presupposed by the Tarskian account of logical truth. Indeed, for reasons that are not entirely unrelated to the Tarskian explanation of "Mars is either dry or not dry", it follows that the universal generalization will also be logically true: "For all objects x, x is either dry or not dry". This is a universal logical truth about *objects*. So one might also say, in some loose sense (not to be made rigorous), that logical truth is "about" the object structure of the world.

3. Platonism

We can now finally return to the question of ontology. As we've seen in the first section, at least one of the reasons why logical conventionalism fails is that it leaves out the worldly element of logical truth. A logical truth about Mars, such as "Mars is either dry or not dry", makes existential commitment to Mars; so Mars must exist for this to be true. This follows from the metaontological idea that names are ontologically committing. But as we saw in the Tarskian account, this sentence's *status* as a logical truth does not really depend on the meaning of the name "Mars". What makes it a *logical* truth is the contributions from "or" and "not". Either way, what makes it *true* (simpliciter) is that *Mars* is either dry or not dry.

So what must exist for there to be logical truths? Well, any logical truth that mentions Mars presupposes Mars's existence; and any logical truth that mentions Jupiter presupposes Jupiter's existence; and so on. So the full range of logical truths that we *actually* have presuppose the existence of all the things that we'd pretheoretically take to inhabit the world. People, planets, tables, chairs, etc. must exist to witness the logical truths about them. But the *general* phenomenon of logical truth does not depend on

any specific object; there just has to be *some objects or other* to serve as objects of quantification, and there will be logical truths.[9] So, to summarize, the actual full range of logical truths presupposes actual objects, but the phenomenon of logical truth itself does not presuppose any object in particular. This might sound paradoxical, but it isn't really.

The main question that now confronts us is whether this is sufficient, or whether there must be some extra, otherworldly things to ground the phenomenon of logical truth. This is the question of *platonism*, since platonism is the view that brings in a whole realm of abstract truth-makers and truth-bearers to explain logical truth.

The best way I can think of to approach this question is to run through a list of other entities that are purported to explain logical truth. To this end, I'll consider propositions, facts, non-Humean laws, and sui generis normative properties.

## 3.1 Propositions

So far I've been treating sentences as the bearers of logical truth and consequence. And this is for good cause: we reason in our vernacular languages. In a formal logical theory, the logician creates artificial, mathematically-defined languages to model logical truth and consequence. But what they're modelling is informal logical reasoning in ordinary language.

But at this point, many will protest that *sentences* (i.e. language-bound vehicles of communication) are not the primary bearers of truth or falsity. Rather, it is the *propositions* that sentences express. Propositions are the information that are encoded by sentences. Since they are the primary bearers of truth and falsity, they are (presumably) the primary objects of truth preservation and logical truth. Hence, they are (allegedly) the prime subject matter of logic.

It is somewhat difficult to comprehensively evaluate this doctrine because it only gives the illusion of a unified theory. In reality, the propositions of propositionalist theory are defined by the theoretical roles that they are meant to serve, and everything else about them is left for the individual to argue about. For this reason, we don't really get a single propositionalist theory to evaluate; we get a myriad of competing ideas that fall under a single heading. Nonetheless, they all have in common that propositions are supposed to be: (i) whatever is the primary bearers of truth and falsity, (ii) the meanings of sentences in public language, (iii) the objects of the so-called propositional attitudes, and (iv) the bearers of modal attributes. Following this, some (namely, Frege) say that propositions are Platonic objects that reside in a third realm; some say that propositions are constellations of worldly objects and universals (Russell before WWI); and some say that propositions are types of cognitive acts performed by humans (Soames).

I will not argue here that propositionalism is false.[10] For all that I will say, propositions may exist and aptly explain mental attitudes and meaning. I will, however, argue that propositions are not apt to explain *logical* truth and consequence, in particular. Specifically *logical* truth is better explained as a feature of language-bound vehicles of communications—i.e. sentences. I have three arguments for this point.

First, the propositionalist theory appears to misclassify cases of logical truth. It is uncontroversial

---

[9] The Tarskian approach to logical truth and consequence presupposes *non-empty* domains of objects.

[10] Actually, I don't really have a settled view on the matter. I do think, though, that much of what propositions are called to explained ought to be explained by relatively-language like systems of representation. See e.g. the appendix to Chapter 1: The Building Blocks of Truth.

that "Hesperus is Hesperus" is a logical truth, whereas "Hesperus is Phosphorus" is not. The sententialist view of logical truth has no trouble explaining this; it's simply because "Hesperus" and "Hesperus" are the same name, whereas "Hesperus" and "Phosphorus" are not (notwithstanding their common referent). But on this point, the propositionalist theory faces trouble. According to one highly compelling view of propositions—i.e. Russell's—"Hesperus is Phosphorus" and "Hesperus is Hesperus" *express the same proposition*. That is because "Hesperus" and "Phosphorus" refer to the same worldly object, and so the same object composes the proposition expressed by each sentence. So if logical truth was a property of this proposition, then it cannot distinguish between "Hesperus is Hesperus" and "Hesperus is Phosphorus". The upshot is that logical truth is better attributed to sentences than Russellian propositions.

A critic will be sure to point out that this is only a knock-down argument against attributing logical truth to *Russellian* propositions. But there are other theories of propositions that distinguish between "Hesperus is Hesperus" and "Hesperus is Phosphorus". Frege, as a prime example, says that propositions are composed of "senses" rather than objects themselves. ("Sense" is a theoretical term. It's a long story. Let's not get into it here.) And since "Hesperus" arguably encodes a different sense than "Phosphorus", it follows that "Hesperus is Hesperus" expresses a different proposition than "Hesperus is Phosphorus". So one proposition can be logically true while the other isn't.

If my critic wishes to invoke Fregean propositions to explain logical truth, then all I can do here is remind them that they are wading into muddy waters. Over the years, a long list of difficulties with the Fregean position has accumulated.[11] I've written up my own summary of these developments in "Fregeanism vs The Theory of Direct Reference". But for our purposes, I should mention just one of them. Fregean senses are postulated to explain two distinct phenomena, the joint explanation of which can give rise to tension: namely, they're supposed to explain *shared meanings in public languages* and *individual rationality*. Consider Kripke's case of a person who, on two separate occasions, has been told that "Paderewski is a politician" and "Paderewski is a musician", but fails to realize that the name was meant to refer to the same person. Now consider the inference: "Paderewski is a politician; Paderewski is a musician; therefore, there is at least one person who is both a politician and a musician". Is this inference logical, according to Frege's account of propositions? Well, insofar as Fregean senses are supposed to encode the public meanings of ordinary expressions (to explain communication between people), each occurrence of "Paderewski" should have the same sense. In that case, their inference would count as logical. But insofar as senses are supposed to explain individual rationality, and given that the inference shouldn't count as logical for our benighted subject, the two occurrences of "Paderewski" should be assigned different senses. But that pulls away from the idea that senses encode common meanings between distinct occurrences of expressions (within a common language, or between distinct cultural languages).

One obvious thing to say here is that this inference shouldn't count as logical *within the individual idiolect* of the benighted subject who's ignorant of "Paderewski" coreferring to "Paderewski". It's as if in *their* idiolect, the two occurrences of "Paderewski" fail to coordinate as occurrences of the same name. But if *this* is how we explain the illogicality of the inference, then we're back to treating logicality as a property of linguistic vehicles, not propositions.

My second argument against the propositionalist account of logical truth and consequence harkens back to my discussion of the subject matter of logic. Basically, the sentential construal of logical consequence comports with the Tarskian account, whereas the propositionalist account doesn't. Moreover,

---

[11] Examples: Putnam's "Meaning of 'Meaning'", Kripke's *Naming and Necessity,* Kripke's "A Puzzle about Belief", Perry's "The Problem with the Essential Indexical", Kaplan's "On Demonstratives".

as I've already argued, I believe that the Tarskian account gives the best analysis of logical truth and consequence.

The Tarskian account of logical truth and consequence hinges on reinterpreting the non-logical names and predicates that comprise the relevant truth-bearers. For instance, we say that "Mars is either dry or not dry" is a logical truth because, no matter how "Mars" and "dry" get interpreted, the sentence will always come out as true. But the natural setting for talking about re-interpreting semantically-endowed items is *language*, not propositions. Whatever propositions are, they *are* meanings. The proposition <Mars is either dry or not dry> will be composed of <Mars> (i.e. the meaning of "Mars"), <dry> (i.e. the meaning of dry), and so on. And it just doesn't make any sense to speak of re-interpreting the meaning of <Mars> to have a different meaning. Whatever serves as the meaning of "Mars", it means *Mars* essentially.

This explains a cryptic remark I made earlier. When we were discussing the modal account of logical truth and consequence, I mentioned that it paired well with the idea that propositions are the primary arena of logical truth and consequence. Well, we can now see why this is. Since the propositionalist cannot straightforwardly adopt the Tarskian account of logical truth and consequence, they must seek instead an account that appeals to features that are compatible with propositions. The natural idea is to appeal to modality. On this account, a set of propositions $\Sigma$ entails proposition *p* if, and only if, it is *necessarily* the case that p is true if all of the members of $\Sigma$ are true.

However, as I argued earlier, a contemporary understanding of logicality and modality drives a sharp wedge between them. Logical consequence is not necessary consequence, and logical truth is not necessary truth. I've already given this argument, so I won't reiterate it again. But now we see how this point suggests a sententialist understanding of the bearers of logical truth and consequence.

My final argument against the propositionalist approach stems from how we express generalizations of logical truths. At least in formal practice, we express the generalizations using sentential schemas. Thus to capture the common logical flavour of "Mars is either dry or not dry", "Jupiter is either hot or not hot", "Graham is either asleep or not asleep", and so on, we produce the schema: $\alpha$ *is either* $\varphi$ *or not* $\varphi$. We can also express it using sentential schematic variables, in that case, it becomes *either S or not S*.

The idea behind schemas is that they are themselves meaningless shells of sentences, but they provide recipes for logical truth through the uniform substitution of actual meaningful sentences, names, predicates, and so on. This way of capturing logical generalizations is inherently sentential. Indeed, if one insists on only ever using objectual quantifiers to express generality, then the schemas become universal generalizations over sentences.

The propositionalist, on the other hand, will want to express the generality of logical laws as true generalizations over their special kind of entity, the propositions. How might they do this? Consider the logical law that's captured by the schema ⌜P or not P⌝, and instantiated by such truths as *either it's snowing or it's not snowing, either Mars is dry or Mars is not dry*. The propositionalist might claim that the law is spelled out by:

- For all propositions *P, P or not P*.

The problem, however, is that this is ungrammatical. "P" is used here both as an objectual variable and as a variable that occupies a sentential position. As a variable ranging over objects, it makes no sense to say "P or not P", since "or" and "not" are sentential/propositional connectives. To make this generalization

work, we're going to have to replace "or" and "not" with expressions for object-level properties and relations.

This point is often obscured by the fact that many people who speak informally about the "laws of logic" get really sloppy and speak only of laws that ascribe object-level properties, like truth and falsity, to objects. Thus they'll say that the law of the excluded middle is that every proposition is either true or false, and this makes good sense.[12] But there's *so much more* to logic than just the laws governing truth and falsity. There's also the truths that stem from the meanings of negation, conjunction, disjunction, the conditional, first-order quantifiers, modal operators, and so on.[13]

So how may we rewrite the above idea so that we grammatically express a law that quantifies over propositions? One idea is to take propositional schema, like ⌜P or not P⌝ as a recipe for cooking up a description of a complex proposition. In that case, the above law should be rewritten as:

- For all propositions *P,* the proposition that is the disjunction of *P* and the negation of *P* is true.

Now we have something that's perfectly grammatical.

The problem, now, is that this law doesn't so straightforwardly tell us *which* propositions are the logically true ones. Suppose that I told you that "Every proposition that Spock is certain of is logically true". What I have said may very well be true, but it wouldn't tell you which are the logical truths until you had the means of matching the description "proposition that Spock is certain of" with specific propositions. Likewise, the above law doesn't tell you which propositions are logically true until you have the means of matching the description "the proposition that is the disjunction of *P* and the negation of *P*", for any given *P*, with a certain proposition. In other words, we need something like a "denominalization" law (a propositional analogue of a disquotation principle): a law that takes us from a description of a complex proposition to a statement of what that proposition says. To put it bluntly, suppose that we have a law instantiated by:

(1) The proposition that is the disjunction of <It is raining> and the negation of <It is raining> is true.

We still need more principles to take us from (1) to (2):

(2) <Either it is raining or it is not raining> is true

Now, you might think that in this case, unlike the Spock case, that this is easy to do. The proposition that matches the description "the negation of P" is the one expressed by "It is not the case that P", and the proposition that matches the description "the disjunction of P and the negation of P" is the one that is expressed by "P or not P". So if you are at all familiar with the technical terms "the negation of" and "the disjunction of", then you have a grasp of how to describe the propositions that result from logical operations, and so you can easily work backwards from a description of a proposition to what the

---

[12] There are two versions of the law of the excluded middle. There's the one that says that every sentence of the form ⌜P or not P⌝ is a logical truth, and there's the one that says that for every proposition P, either P is true or it is false. These must be kept separate, as they can come apart in non-classical logics.

[13] Don't say that we can rewrite the above law as "for all propositions *P, either P is true or its negation is true"*. The problem with this is that it *uses* "or" to describe a law governing truth and negation; it does not describe any law *governing* disjunction.

proposition says, or how it was concocted through the use of logical operations. So in this case, we do have the means of matching the description of a proposition with what the proposition says.

Nonetheless, this still shows us that there is a gap between the law that *describes* the complex proposition that results from logical operations (and ascribes it truth), and the laws that govern the *use* of logical operators to form complex, logically true propositions. To bridge the gap between them (and thus, to understand the description of propositions), we are implicitly relying on schematic principles of the form:

(D) The proposition that is the disjunction of <A> and <B> says that A or B, and is expressed by "A or B"

(Neg) The proposition that is the negation of <A> says that ~A, and is expressed by "It is not the case that A".

Once again, what we have here are *schemas*, not fully interpreted laws about propositions. They cannot be understood as quantifying over propositions because the second occurrence of "A" and "B" in (D) and the second occurrence of "A" in (Neg) are not in object position. The only way to understand these principles is through the uniform substitution of sentences in the "A" and "B" positions.

If I'm right about this, then the natural way that we understand the *generality* of logical laws is through the substitution of sentences into uninterpreted schemas, and not through quantification over propositions.[14]

One might try to resist this conclusion by positing propositional *structure* that mirrors the syntax of language. After all, we understand the significance of the sentential schema ⌜P or not P⌝ because we understand the syntax, and this allows us to think of the "P" position as a kind of hole that can be filled by any sentence whatsoever, resulting in a grammatically complex sentence. So perhaps the propositionalist will claim that, besides the propositions, there are also propositional structures. There are abstract entities that can be represented by <__ or not __> that take propositions as arguments and yield complex propositions as a result. (Example: <snow is white> inputed into <__ or not __> yields the proposition <snow is white or snow is not white>.)

Upon positing these abstract propositional structures, the propositionalist might recast the logical laws in terms of them. For instance:

● For all propositions P, P inputed into the propositional structure <__ or not __> yields a truth.

This too is grammatical. Although I can't help but point out that these propositional structures are mysterious, and that our only means of understanding them is by projecting our understanding of sentential syntax. In other words, our understanding of this principle is undoubtedly parasitic on our understanding of the logical form of *language*.

Regardless, there is a pretty significant—and often overlooked—reason to be wary of a realm of

---

[14] I can foresee someone trying to resist this argument by piling on further descriptions to pin down the identity of complex propositions. For instance, one can rightly say that, for all *p*, the negation of *p* is the proposition that results from inputting *p* to a negation operator, and it is true if, and only if, p is false. All of this is fine, but still, to go from a description of the proposition to what the proposition says, at some point we'll need to appeal to substitution schemas.

abstract *structured* propositions that are required for this principle. The problem is that the principles that characterize structured propositions also entail a paradox—a contradiction. This is no different from how naive set theory entails a contradiction through the construction of *the set of all sets that do not contain themselves*, so we know that the sets of naive set theory do not exist. This is called the Russell-Myhill Paradox (also Russell Appendix B paradox). I'll link the reader to an explanation: https://iep.utm.edu/par-rusm/. Basically, consider the class *w* of propositions that state that all the members of some class of propositions, of which they are not a member, are true. Since this itself is a class of propositions, we can consider the proposition *p* which states that *all of w's members are true.* The question now is whether *p* is in *w*. By using an inevitable assumption about propositional structure, we can derive the result that *p* is in *w* if, and only if, *p* is not in *w*.

Needless to say, the charge of contradiction should not be taken lightly. And there would be an odd irony, in this dialectical context, of invoking structured propositions to explain *logical* laws while shirking off a literal logical contradiction. I'm not going to say that this problem can't be overcome (perhaps one could retreat to unstructured propositions, or place arbitrary limits on how many propositions there are). But add this to the list of the other three or four problems I've adduced for the propositionalist account, and you see why I favour sententialism.

3.2 Abstract sentence types

Given that I've said that logical truth pertains to sentences rather than propositions, I'd imagine a critic pointing out that I haven't thereby avoided ontological commitment to abstract entities. Sentences, after all, are a kind of abstract object.

The critic would be quite right, but misunderstand my intentions. None of my complaints against the propositionalist account of logical truth concern the abstractness of propositions. Rather, most of my arguments concern the fact that the meanings of sentences are ill-suited to do double-duty as the vehicles of logical truth, given the Tarskian conception.

So I fully accept that the meta-logical account of logical truth entails the existence of a kind of abstract entity: sentences. (The logical truths themselves will also commit us to this kind of entity, since there are logical truths about sentences. Case in point: "'Mars is red' is either grammatical or it's not grammatical"; this is a logical truth, and it commits us to sentences.) So we do not avoid the existence of abstract objects altogether.

Now that I've introduced this second kind of ontological commitment, I should say a few words about what I take these truth-bearers to be. First, we must distinguish between concrete linguistic tokens and abstract sentences. Concrete tokens are produced by particular acts of speech; they occur at a particular time, in a particular context, and so on. Concrete tokens will have a physical manifestation (marks on a page, sounds in the air), but they are so much more besides. They are also produced by speakers with particular thoughts and intentions, and have particular causal histories, which imbues them with their semantic features (i.e. reference and truth conditions). I've written elsewhere why I regard these to be the primary bearers of semantic significance when it comes to explaining the determination of reference.[15]

But when it comes to *logical* truth and consequence, we need to focus on abstract linguistic *types*, not tokens. This is because a sentence's status as a logical truth, or a sequence of sentences's status as a logical consequence, crucially depends on the *re-occurrence* of one or more linguistic types. Take, for

---

[15] https://grahamsethmoore.wordpress.com/2021/12/31/truth-bearers-i/

example, the inference of "there is at least one person who is both a musician and politician" from "Paderewski is a musician" and "Paderewski is a politician". This sequence constitutes a logical consequence only if the first and second occurrences of "Paderewski" count as re-occurrences of the same abstract word type.[16] Hence the need to invoke abstract linguistic types to account for logic.

Let me say a few words about how I think of the abstract languages that serve as the arena for logical truth. First, they are abstract in the sense that their expressions are multiply-instantiable. Second, they are also abstract in the sense that they're completely characterized by (a) the syntactic rules that define the well-formed formulas (what's to count as a formula of a language), and (b) the assignment of denotations to the simplest expressions of the language, and (c) the semantic rules for determining the truth conditions of complex sentences.[17] Given that they're semantically interpreted, it would be wrong to think of the sentences of these languages as mere physical items; they're not mere words on a page or sounds in the air. But on the other hand, this is still a fairly austere way of thinking of a language, and it is still several steps removed from the living, breathing linguistic practices of human beings. For instance, these languages are set up in a way so that each lexical item bears its denotation essentially. As a result, it makes no sense to ask *how* or *why* "horse" in an abstract language refers to horses. It *does* make sense to ask why a speaker's use of "horse" in their idiolect refers to horses, but the answer to this will presumably advert to psychology (i.e. their and their community's cognitive rapport with horses). It is entirely appropriate for the purposes of *logic* to appeal to theoretically-abstracted languages that ignore these metasemantic details. As Frege and the early Analytic Philosophers were wont to stress, logic is not psychology. As such, we abstract away from the psychological details of language use.

Where do these abstract languages come from? I think it's most plausible to think of them as abstract *artifacts*. *Qua* abstract object, their characterizations abstract away from logically-irrelevant details of human language use. But *qua artifact*, they still are *created* by human linguistic practices. Somehow, through our linguistic conventions, we manage to produce abstract languages.

Some philosophers, stuck in a platonic rut, will object that abstract objects, by their nature, cannot be created because they are acausal. According to them, all abstract objects dwell in a platonic realm, cut off from our activities.

But there are many ways of defining "abstract", and it seems to me to be mere dogma to think that some of the definitions imply the others. I said that abstract languages are characterized by disregarding some (but not all) concrete features of language use and that their expressions are multiply insatiable. How does that imply that they're causally inert? This would need an argument.

Besides, this implication is not very plausible on the face of it. Pre-theoretically (that is, before we were bullied by platonists), we would all agree that symphonies, stories, and laws are artifacts. They were created by people and did not exist for all of time. And yet they are abstract. A symphony is not identical to any one performance, and a story is not identical to any one telling. They abstract away from particular details of their instantiations; e.g. a symphony is not defined to require *this* particular trombone

---

[16] This argument can be resisted. There's space for a view that takes the coordination between token words as a primitive relation that is not explained by the instantiation of the same abstract object. I take it that Kit Fine's 'semantic relationism' is such a view. But I won' t explore it here because I'm not super familiar with it.

[17] If an abstract language is to be suitable for the purpose of logic, it must formally distinguish between semantically-distinct expressions that have a morphologically-identical expression in ordinary language. In other words, "bank", as used for financial institutions, and "bank", as used for the sides of rivers, must count as "distinct words" in the characterization of such a language. There is a hugely interesting question of how expression types are individuated in these abstract languages, but however it's answered, for logical purposes, it must be constrained by semantics.

player, or be played on *this* particular wood instrument.

Anyway, there are a lot of *bad* arguments against the idea that logical truth is a property of sentences. In trying to be as complete as I can, I'll mention three more.

*Objection 1. If sentences depend, for their existence, on being spoken or thought by actual humans, then you cannot say that there are logical truths that have never been spoken or thought. But there are logical truths that have never been spoken or thought. So logical truth cannot be a property of sentences.*

*Response*. This argument relies on a false premise: that sentences depend, for their existence, on being spoken or thought by actual human beings. I mean, that's just not true of languages conceived of as abstract languages. The sentences of such languages are defined by the recursive rules that generate the well-formed formulas. (Example: consider the rule "if A is a sentence and B is a sentence, then ⌜A and B⌝ is a sentence". This rule alone would guarantee the existence of arbitrarily long conjunctions, most of which are too long for finite humans to speak.) Most well-formed sentences of English (for example) are such that they've never been spoken by any human being. Many of them are logical truths.[18]

*Objection 2. If sentences are artifactual, then they are conventional and mind-dependent. But if logical truth is a property of sentences that are conventional and mind-dependent, then logical truth is conventional and mind-dependent. But logical truth is not conventional or mind-dependent. We cannot make it true that it's raining and it's not raining. We cannot decide that it's false that Mars is either dry or not dry.*

*Response*. This objection commits a fairly egregious vehicle-content conflation fallacy. In short, the second premise is false. A sentence is a mind-dependent object, but whether its truth condition obtains, given the meaning that it has, can be a fully objective thing. Case in point: the sentence "Jupiter has 92 moons" is an artifactual, mind-dependent creation. But given that "Jupiter" refers to Jupiter (a conventional, mind-dependent fact) and given that "has 92 moons" applies to things with 92 moons (another conventional, mind-dependent fact), and given that the sentence's structure represents property attribution (a third conventional, mind-dependent fact), it follows that it is true if, and only if, Jupiter has 92 moons. And the latter condition is patently not up for humans to decide. Well, the same goes for logical truth. Given that "Mars", "or", "not", and "dry" mean what they do, it's not up for humans to decide whether the truth condition of "Mars is either dry or not dry" obtains.

On a similar token, a map is quite clearly an artifactual creation. But it does not follow that the artificer's whim determines whether their created map is accurate. If I create a map of Europe that depicts Paris as north of London, then my map is objectively wrong. Regardless of the map's status as mind-dependent, it's a mind-independent fact that London is north of Paris.

*Objection 3. If sentences are artifactual, then they exist contingently. But logical truths are necessary.[19] Hence, whatever bears logical truth cannot be contingent; the truth-bearers must exist necessarily.*

---

[18] What about logical truths concerning faraway planets for which we have no name? Well, we have demonstrative expressions—"this" and "that planet"—that can be assigned those far away planets as referent on some interpretation. In this way, the English language—*qua abstract language*—has the means for capturing those logical truths.

[19] Recall that in section 2 I argued against this premise. "It's raining if, and only if, it's actually raining" is a logical truth that is not necessarily true. But we'll leave this point aside now.

*Response.* I reply that there's no incoherence in attributing *truth-by-necessity* to a contingent truth-bearer. Indeed, modally-graded truth can be cashed out as a semantic property of any truth-bearer that represents things as being a certain way—regardless of whether the truth-bearer exists necessarily. It is the property that a truth-bearer has if how it represents things as being is a way that they *must* be. For example, to say that "Water is $H_2O$" is true-by-necessity is just to say that "Water is $H_2O$" represents water is $H_2O$, and that water is necessarily $H_2O$. But notice that this account doesn't require the sentence "Water is $H_2O$" to itself be a necessarily existing entity. Well, the same goes for logical truth. "Either it's raining or it's not raining" is a logical truth that is also true-by-necessity. This is because it represents things such that either it's raining or it's not raining, and things *could not fail to be such* that it's either raining or it's not raining. We thus get an account of the sentence's truth-by-necessity that makes no mention of the sentence itself existing necessarily.

I know very well that the main reason that people demur from this point is because they have an alternative account of necessary truth in mind. They think that a truth-bearer P is necessarily true if and only if P exists in every possible world (as a worldly denizen) and is true in every world that it exists. I have written a rather long critique of this idea [here](https://grahamsethmoore.wordpress.com/2020/02/11/in-defence-of-the-contingent-existence-of-propositions/).[20] So I won't bother telling that whole story again; an interested reader should just go to the other post. Suffice to say that this alternative account of necessary truth is neither inevitable nor attractive. It is not inevitable because I just gave you a better alternative in the previous paragraph (and have given more details elsewhere). It is not attractive because it appears to presume that a proposition must exist *in* a possible world in order to accurately represent it. But if a proposition is *representational*, then I just don't see this idea having any traction. After all, a map doesn't have to exist *inside* a country in order to represent it; nor does a sentence have to exist *at a time* (e.g. 3,000 years ago) in order to truthfully say something *about that time*. Nobody is baffled by the fact that the English sentence "Plato founded the Academy" is true about an event that occurred in 389 BC, despite the *English sentence* not yet existing in 389 BC. Likewise, nobody should be baffled by the fact that the English sentence "the British Isles were never inhabited" is true of a non-actual, counterfactual world, despite the *English sentence* not existing in that possible world.

## 3.3 Non-Humean Laws and Facts

From now on, I'm going to take it as granted that logical truth is a property of abstract artifactual sentences, rather than propositions. This brings the metaphysics of logic a significant step down to Earth, and out of the platonic third-realm. However, it does not yet steer us clear of all of the ontological questions that one might raise about logic. There is another kind of metaphysical-maximalist who will think that we've neglected the most important entity of all: the laws of logic themselves!

According to this person, it was always a mistake to identify the laws of logic with propositions. We should never have identified the law of the excluded middle with the proposition <for all P, <P or not P> is true>. That is because propositions are *representational*, and they are made true by the facts that they represent. The proposition <for all P, <P or not P> is true> *represents* the law of the excluded middle, and is *made true* by the law of the excluded middle, but it is not the law itself. We should never mistake a representation of something with the thing itself.

If the laws of logic are distinct from the propositions, then what are they? At this point, I can see the metaphysical-maximalist attempting one of two strategies to motivate reifying the laws of logic. I say

---

'reifying' because the end result of their efforts is to identify each law of logic with a specific *entity*. They want to say that the laws of logic are specific denizens of the world, and as such, incur robust ontological commitments. As we'll see, there's something a bit fishy about all of this.

The two strategies they may attempt are: (i) invoking non-Humean laws as explanatory posits, and (ii) identifying the laws of logic with certain *facts*. We'll discuss each of these in turn.

3.3.1 Non-Humean Laws

First, one might think that laws of logic must be posited as a kind of *sui generis* abstract entity in order to explain a range of *logical regularities*. Consider this list of logical truths:

> I. Either it is snowing or it is not snowing
> II. Either it is sunny or it is not sunny
> III. Either $1 + 1 = 2$ or it is not the case that $1 + 1 = 2$.
> …

And so on. Clearly, all of the instances of this list have something in common: they are all instances of the form "P or not P", they are all true, and they all appear to have a common logical explanation for their truth. Informally, they are all instances of the law of the excluded middle.

Noticing these similarities, a reifier might make the following argument: *It is not enough to explain the truth of propositions I, II, III, … severally; they all must have a common explanation. Indeed, they are all made true by a single law: the law of the excluded middle. Hence this law must exist in order to serve as the common explanation for I, II, III,.... But if this law exists, then it is not identical to any one of the I, II, III,...; it must be something 'over and above'. It must be an abstract entity that explains all of these truths.*

This argument is typical of a *non-Humean* view of laws. According to this view, laws must exist, as sui generis entities, over above their instances, to explain their common occurence. This contrasts with Hume's view of laws of nature, according to which a law is simply a true universal generalization—no entity needs to exist in addition to what's presupposed by the instances.

Consider the fact that all ravens are black (well, ignoring the albino ones and the white-necked ones in Tanzania). Since this is explained by biology, you might even call it a law. Let's call it the "law of the blackness of ravens". Ontologically speaking, what needs to exist in order to explain this law? A Humean would say that the law of the blackness of ravens consists of nothing more than the following facts:

> 1. Raven1 is black
> 2. Raven2 is black
> 3. Raven3 is black
> …
> N. Non-black thing-1 is a non-raven
> N+1. Non-black thing-2 is a non-raven
> N+3. Non-black thing-3 is a non-raven
> …

And so on, for every raven that there is and every non-black thing that there is.

And so as far as our *ontology* is concerned, the entities needed to witness the law are simply all of the ravens and non-black things.

The non-Humean, by contrast, would want to see another entity: the law itself. We can imagine a child being told of the law of the blackness of ravens, and wanting to see a demonstration. In order to prove to them this law, we parade before their eyes every single raven in existence and every single non-black thing. After seeing every raven and every non-black thing march by, we can imagine the child protesting; "I saw all of the ravens, and I saw that they were all black. I saw all of the non-black things, and I saw that none of them is a raven. But what I never saw was *the law of the blackness of ravens* march by. I can't know that all ravens are black because the *law of the blackness of ravens* never participated in the parade."

I dare say that this child has committed something like the category mistake discussed by Gilbert Ryle. In one famous episode, Ryle used the following parable to illustrate the concept. A child is watching a parade of soldiers. He sees the battalions march by, and then the batteries, and then the squadrons. After these, the child exclaims "I've seen the battalions, the batteries, and the squadrons, but when is the *parade* coming?" As Ryle points out, the child has made a mistake. The child confusedly thought that "a parade" refers to one more division of soldiers, when in fact it refers to the whole thing.

According to the Humean, the non-Humean is making something like Ryle's category mistake in their ontological expectations. For them, the law that all ravens are black does not demand any extra entity over and above all of the ravens and non-black things. There is no additional thing—the law—that needs to march down as part of the parade. For the Humean, speaking of "the law of the blackness of ravens" is just a roundabout way of saying that all ravens are black.

But for the non-Humean, this is unsatisfying. For them, a *real* law cannot be just a random assortment of instances of a true universal generalization. There is no *law*—they would say—that all U.S. presidents have been men: it's just a true universal generalization that isn't a law. What separates a mere true generalization from a lawful generalization is that there is a law that explains the latter.[21]

When it comes to the blackness of ravens, I am squarely on Hume's side (and I hope the reader will agree). We don't need, in addition, all of the ravens and non-black things, any other entity to explain why all ravens are black. To be clear, this is *only* a point about the ontology needed to account for general facts. The Humean and the non-Humean disagree over what exists. The non-Humean need not say that there is no common explanation for all of the particular facts that entail a general fact. Indeed, for the ravens, there *is a* common explanation as to why raven1 is black, raven2 is black, and so on. Namely, each raven comes from the same biological order, possesses similar DNA, and this DNA encodes for black feathers.

Now let's return to logic. The question before us is strictly one of ontology. Each side takes it as granted that (I) Either it is snowing or it is not snowing, (II) Either it is sunny or it is not sunny (III) Either $1 + 1 = 2$ or it is not the case that $1 + 1 = 2$…, and so on, for each statement of the form "P or not P". Where they disagree is over the ontology needed to explain these facts. The Humean claims that our ontological commitments need not extend past whatever is required by I, II, III,... taken severally. The non-Humean thinks we need something more; they think that in addition to the truths I, II, III,... there

---

[21] In this case, the Humean has a response: genuine laws can be sorted out from accidentally true universal generalizations by appealing to what's *necessary* and what's *possible,* without introducing new entities. Since it's possible for a woman to be president, it follows that "all presidents are men" doesn't describe a law.

must be another abstract entity—the law.

The best argument against the non-Humean is that, in this case, we can explain the common truth of I, II, III,... without invoking a *sui generis* law. I, II, III, … and every other substitution instance of "P or not P" are each complex sentences formed by a common sequence of operations. The sequence starts with a sentence, *S*, forms a negation, *not-S*, and then forms the disjunction, *S or not-S*. Moreover, given that "or" contributes the truth function associated with disjunction, and "not" contributes the truth function associated with negation, it's easy to explain why *any* sentence formed by this sequence of operations will yield truth for any truth value of S. In other words, by invoking the common *form* of I, II, III, …, and the semantical contributions of "or" and "not", we get a fully general explanation of the truth of every instance. And nowhere in this explanation did we resort to saying "because there is this law that makes it true."

Indeed, the explanation in terms of syntax and semantics is actually much more powerful than what's provided by the non-Humean reifier. That is because it is an application of much more general semantical principles that can explain logical truths of any type. By appealing to syntax and the semantic contributions of "or", "and", "not", and "if, then", we can recursively explain the logical truth of virtually every propositional tautology, just by using truth tables. This is a much better, more satisfying, explanation than reifying a logical law for every form of PL tautology. The more that we can explain with a smaller set of principles, the better our explanations are. So the non-Humean *suis generis* law is not only unnecessary, but it is also inferior as an explanation.

3.3.2 Facts

As I've portrayed them, the previous kind of reifier makes their case as an inference to the best explanation. They say that non-Humean laws *best explain* the regularities in logical truth. (And I reply that they aren't the best explanation; a better one is provided by the semantical properties of the logical connectives.) But there are other kinds of reifiers who care less about explanation and instead launch their position from metaphysical first principles.

This other kind of reifier I have in mind is one that subscribes to a general ontology of *facts*. Roughly speaking, there are two approaches to ontology: fact-based and object-based. An object-based ontology takes as its starting point the existence of objects—quarks, electrons, tables, chairs, people, etc.—and perhaps the properties that they instantiate. A fact-based ontology, on the other hand, conceives of the world as ultimately made up of facts. (Wittgenstein asserted so in the opening of the *Tractatus*.) In such a view, there may also be objects and properties, but a simple inventory of the objects of the world will fall short of everything that there is. Because in addition to (for example) the cat and the mat, there must also exist the *fact* that the cat is on the mat.

This division over object-based and fact-based ontologies isn't really a dispute in the metaphysics of *logic* per se; rather, it's a much wider dispute over general ontology. It's most likely to arise as an in-house debate between correspondence theorists over the nature of correspondence truth. Some correspondence theorists (i.e. early Wittgenstein) would insist that facts exist to serve as the correspondents of true propositions. Other correspondence theorists disagree and say that truths correspond to objects and their properties.[22] Regardless, I can foresee how this debate could be relevant to the present issue. One kind of reifer will argue that the laws of logic *exist*—they are denizens of reality—and they exist in the form of facts. Again, this reifer will think it's a mistake to identify the laws

---

[22] https://grahamsethmoore.wordpress.com/2023/02/14/chapter-1-the-building-blocks-of-truth/

of logic with any proposition. They will say that the law of excluded middle is not identical to the proposition <for all P, <P or not P> is true>. Instead, they will say that the law of the excluded middle is the *fact* to which this true proposition corresponds.

(There is a certain irony in this position in that the foremost defender of the fact-based correspondence theory is early Wittgenstein, and yet he explicitly argued that facts only correspond to *non-logical,* contingent propositions. According to him, logic is a matter of the *form* of propositions, not their content.)

Since this kind of reifier appeals to a very broad (and controversial) ontological position, I can't really dispute their specific claim that LEM exists in the form of a fact. Unless you're early Wittgenstein, if you believe that facts exist over and above the objects that populate the world, then presumably you'll think that there's a fact that for all propositions P, <P or not P> is true.

For my part, the best I can do is argue that the whole fact-based ontology is superfluous and that an object-based ontology is all we need. To make this case in full, I would need to flesh out a correspondence theory of truth that takes objects as the only ontological posit. Lucky for me, I happen to have already done that (see [The Building Blocks of Truth](#)). The basic idea is to build one's correspondence theory on top of a theory of reference for basic lexical items, like names and predicates, and from then on, define truth-by-correspondence using a short list of recursive principles to generate truth conditions for sentences of any complexity. (This is an approximation of the view. Hartry Field took this exact line in "Tarski's Theory of Truth", and it needs some adjustments, which I make in my paper.)

Such an approach to truth will give a full explanation of all of the object-language logical truths without invoking any facts. Take the (supposed) logical truth that Mars is either dry or not dry. Given (i) an explanation of how "Mars" refers to Mars and "is dry" applies to dry things, (ii) the semantic contribution of predicating "is dry" upon Mars, and (iii) the truth functions associated with "or" and "not", we get the result that "Mars is either dry or not dry" is true if, and only, Mars is either dry or not dry, which will yield truth no matter the truth value of "Mars is dry". In a similar fashion, we can recursively explain the truth of virtually all first-order logical theorems. And with additional semantic apparatuses, we can extend our reach to other kinds of logical truths (higher-order, modal, temporal, etc.).

Once we have an explanation of all object-language logical truths, it's a short step to expressing generalizations. The standard way to do so is to use schematic variables. For instance, we can say that all sentences of the form ⌜P or not P⌝ are logically true. And as outlined in the previous section, we can offer a common explanation for why all instances of this schema are logical truths. Again, all that needs to exist for these explanations to go through is the objects mentioned by the various operand sentences (the 'P's). There's no need to invoke facts as additional ontological posits.

So why prefer an object-based ontology over a fact-based ontology? The short reason is that the facts of the fact-based ontology don't seem to do any crucial work. I agree with Quine when he writes:

> What on the part of true sentences is meant to correspond to what on the part of reality? ... [P]erhaps we settle for a correspondence of whole sentences with *facts:* a sentence is true if it reports a fact. But here again we have fabricated substance for an empty doctrine. The world is full of things, variously related, but what, in addition to all that, are facts? They are projected from true sentences for the sake of correspondence. (Quine 1987, pg. 213)

And if the Tarskian conception of logical truth and inference is correct, then we can explain *logical* truth, in particular, just in terms of objects and their properties, and without mentioning facts.

If the world consists of objects variously related, then how come we ever speak of facts? A critic might point out we often mention facts in other contexts besides theorizing about truth.

In my view, "fact"-discourse plays a role analogous to what deflationists claim that "truth"-discourse plays. That is, according to my theory, to say that "it is a fact that P" or "there is a fact that P" is really just a fancy way of (more or less) saying "P". To say that *there is a fact* that snow is white is really just a fancy way of saying that *snow is white.*

But if that's the case, then why would we ever use the stilted, longwinded "there is a fact that P" instead of simply saying "P"? Once again borrowing from the deflationists, my answer is that it is sometimes convenient to speak of facts in order to say many things at once. It would take me forever to list off my intrinsic attributes (*Graham is human, Graham is 5'10", etc.*), but luckily I can take a shortcut and speak of "all of the intrinsic facts about me". This gives us a good reason to introduce vocabulary that apparently refers to facts. However, as long as we understand it as a mere convenience rather than revealing a serious ontological commitment, we can hew to an object-based ontological picture.

Admittedly, this is rather quick. It would take me much longer to satisfactorily lay out the commitments of my deflationist theory of facts. I've already undertaken the first few steps [here](https://grahamsethmoore.wordpress.com/2022/01/08/facts-as-fictions-or-the-deflationary-theory-of-facts/).[23] But let's not get sidetracked by the metaphysics of facts. The main point has already been made: a fact-based ontology isn't necessary for explaining logic, and if my theory is right, then it isn't necessary *simpliciter.*

3.4 Normativity

There is one more aspect of logic that deserves our consideration (I may very well have saved the best for last). For many people, logic is special because it is *normative*. Logic doesn't merely tell us what *is* true; it also tells us *what we ought to believe* and *how we ought to reason*. For them, the whole point of outlining a system of logic is to uncover the rules for reasoning and thought.

If this is right, then it isn't enough to explain how the logical theorems describe truths. A proper metaphysics of logic must also explain the *normative authority* that logical laws exert over our thought. We might put this point by saying that logic gives rise to a certain distinctive normative property—let's call it the "logical ought" or $ought_L$—that figures into the norms that guide our beliefs. Among the norms of belief, there will be distinctly logical ones like:

- One $ought_L$ not believe contradictions
- If one knows <P> and <P→Q>, then one $ought_L$ to believe <Q>.

And the job of a metaphysics of logic is to explain this distinctive normative property.

I know that a lot of people are gripped by this idea. And so they're going to have to excuse me for being bluntly dismissive: but logic just isn't normative in the way that they imagine.

Even on the face of it, logic, as a discipline, is concerned with certain kinds of *truths* and *truth-preservation* across arguments. It is concerned with the sentences that are true in virtue of their form and the semantic contributions of the logical constants; it is also concerned with the arguments that preserve truth in virtue of form the premises and conclusion, and the semantic contributions of the logical constants. So far, all of this is descriptive; it tells us *which* sentences are true and *which* arguments are truth-preserving. Indeed, logical theories hardly ever invoke normative terms. (The only uncontroversial exceptions to this rule are the logical theories *about* normative terms—i.e. deontic logics.)

---

[23] https://grahamsethmoore.wordpress.com/2022/01/08/facts-as-fictions-or-the-deflationary-theory-of-facts/

So if one wants to claim that logic is normative *despite* not appearing so, then one must argue that the normativity gets smuggled in through the back door. One must say that even though logical theorems *appear* to be descriptive, they're *really* normative under the surface.

To be sure, there have been a number of philosophers who have tried to make this case. To respond to them, I can do no better than Gillian Russell in "Logic isn't normative" (2020).[24] I refer the reader to this paper to make a more elaborated case for the present point. (Since this is a large topic, I can't accomplish anything more than a brief summary and a gesture at more detailed resources.)

When I say that logic isn't normative, I mean that logic does not give rise to a *distinctive* kind of norm (the "logical ought" or "*ought$_L$*"). To be sure, I'm not saying that there is no norm against believing contradictions, or that there's no norm of completing a modus ponens from known premises. But this normativity can be entirely traced back to very general norms that guide belief *in general*. Arguably, there is a norm that says we ought to believe that which is true (or that which we are in a position to know). This is a very general norm; it concerns *all belief*; it is not specific to logic *per se*. But nevertheless it explains why we shouldn't believe contradictions: because contradictions are never true! On that token, we also shouldn't believe that the moon is made of cheese. The reason why we shouldn't believe this is not because it's a violation of any law of logic, but simply because it's false.

Similarly, because there is a general standing norm in favour of believing the truth, it follows that one should complete modus ponens arguments from known premises. If you know that P and you know that P implies Q, then you ought to believe that Q. Again, this is not because modus ponens is an *intrinsically* norm-generating form of inference; rather, it's because the validity of the argument guarantees the truth of conclusion, and because you should generally believe true things (that you're in a position to know).

If this is what the supposed normativity of logic amounts to, then logic is no more normative than mathematics, physics, biology, etc. All of these other fields deal with describing truths, and to that extent, are authoritative over what we should believe. If you know that you have two coins in your left pocket and three coins in your right pocket (and you have no other pockets), then you should believe that you have a total of five coins. So mathematical truths, like 2 + 3 = 5, have implications for what you should believe. But nobody would conclude from this that mathematics is a normative discipline. Nobody would demand a metaphysics of mathematics to account for the 'normative' element of mathematics. Rather, we rightly see that mathematics *describes* truths, which, when taken in conjunction with a standing obligation to believe the truth, will have normative consequences. The normativity isn't *intrinsic* to math.

Well, the very same is true of logic. Logic tells us what is (descriptively) true. For instance, it tells us that if <A or B> and <not A> are both true, then <B> will be true. And *of course* this has normative consequences *when taken in conjunction with a general norm that tells us to believe truths and disbelief falsehoods*. But this normativity is not *intrinsic* to logic. It only arises when we combine the truths of logic with an independent norm of belief. And so nobody should think that a metaphysics of logic needs to account for it.

---

[24] https://philpapers.org/rec/RUSLIN