Knowledge is Closed Under Analytic Content

Samuel Z. Elgin

July 2019

Abstract

I defend the claim that knowledge is closed under analytic content: that if an agent S knows that p and q is an analytic part of p, then S knows that q. After identifying the relevant notion of analyticity, I argue that this principle correctly diagnoses challenging cases and that it gives rise to novel and plausible necessary and sufficient conditions for knowledge. I close by arguing that contextualists who maintain that knowledge is closed under classical entailment within—but not between—linguistic contexts are tacitly committed to this principle's truth.

Knowledge of some propositions requires knowledge of others. All those who know that $p \wedge q$ also know that p—as do those who know that $\neg\neg p$. Agents ignorant of the truth of p lack the epistemic resources to know either the conjunction or the double negation. When does this hold? Is it a feature of conjunction and negation in particular, or are they instances of a more general pattern? Why does this hold? What makes it the case that knowledge of some propositions requires knowledge of others?

This is the interpretive question of epistemic closure. The most basic formulation of the closure principle is the following:

> **Naïve Closure:**
> If an agent S knows that p and p entails q, then S knows that q.[1]

Naïve Closure is transparently false. S may fail to realize that p entails q, and perhaps even disbelieve that q. And, surely, if S does not even believe that q, then S does not know that q. We might amend Naïve Closure by requiring S to recognize that p entails q in order to count as knowing that q, but this does not accommodate cases where S believes that q for spurious reasons, rather than because it is entailed by p. And we're off to the races! Philosophers continue to add further conditions, clauses and amendments, as we are wont to do, in order to stave off the perennial threat of counterexample. As accounts continue to expand, closure becomes more complex; what initially appeared to

[1] There are also closure principles concerning justification and evidence, which I set aside for the purposes of this paper. See, however, Kyburg (1961); Wright (1985); Davies (1998).

be an intuitive and straightforward connection between knowledge and entailment quickly becomes an extraordinarily convoluted affair. As with other interpretive debates, it is not obvious that there is a unique correct resolution; perhaps multiple versions of closure obtain. But concerns about an overabundance of uncontroversial interpretations may be premature, as we have yet to uncover one.

Despite the increasing complexity and the arduousness of precise formulation, many continue to find some version of closure appealing.[2] Of course, there are detractors. Dretske (1970, 2005), for example, argues that knowledge involves tracking truth, but the ability to track truth is not closed under entailment. Nozick (1981) maintains that we can account for both the appeal and the falsity of skepticism by denying closure; because closure fails, I know that I have hands, and the fact that I have hands entails that I am not a handless victim of a Cartesian Demon, but I do not know that I am not a handless victim of a Cartesian Demon. However, a great many philosophers continue to find some version of closure attractive—despite our inability to formulate it precisely. It is a truism that we epistemic agents expand our knowledge by recognizing the consequences of what we already know; perhaps closure even commands a kind of Moorean certainty that philosophy is, in principle, unable to undermine.

I maintain that Naïve Closure is perfectly true as it stands—no additional clauses or amendments are required. What is needed, instead, is a refined notion of entailment. Often, philosophers (logicians, mathematicians, etc.) use 'entailment' to refer to a relation that preserves truth; a proposition p entails proposition q just in case the truth of p guarantees the truth of q. However, another conception of entailment concerns the preservation of *meaning*. A proposition p entails proposition q, on this conception, just in case the meaning of q is a part of the meaning of p. If the meaning of 'John is a brother' contains the meaning of 'John is a sibling,' then 'John is a brother' entails that John is a sibling, and if the meaning of 'It is raining' does not contain the meaning of '$1 + 1 = 2$,' then 'It is raining' does not entail that $1 + 1 = 2$—despite the fact that it guarantees the sentence's truth. This is the kind of entailment under which knowledge is closed: the kind between a sentence and its analytic parts. In particular:

> **Analytic Closure:**
> If an agent S knows that p and q is an analytic part of p, then S knows that q.

Any defense of Analytic Closure ought to be preceded by a discussion of the relevant notion of analyticity, so that is where I begin. It is often said that a sentence is analytic just in case its truth-value is determined purely by the meanings of its terms; plausibly, 'All doctors are doctors' and 'All vixens are foxes' are both analytic. In contrast, a sentence is often said to be synthetic just in case its truth-value depends upon more than the meanings of its terms; plausibly 'Some doctors are wealthy' and 'All vixens have tails' are synthetic.

[2] For defenses of closure, see, e.g., Vogel (1990, 2000); Williamson (2000); Hawthorne (2004, 2005).

Kant (1781) conceived of analyticity primarily as a property of judgments, rather than sentences. He provided two, arguably distinct, interpretations of the analytic/synthetic distinction.[3] According to one, a judgment is said to be analytic just in case its negation yields a contradiction; analyticity can be understood in roughly logical terms. Because 'John is such that the law of excluded middle is false' engenders contradiction, 'John is such that the law of excluded middle is true' is analytic. According to the other, a judgment is said to be analytic just in case its predicate is conceptually contained within its subject. If the predicate 'extended in space' is conceptually contained within 'body,' then 'All bodies are extended in space' is analytic. The reason these come apart, some maintain, is that the second interpretation is more fine-grained than the first—it requires, minimally, that the predicates and subjects be relevant to one another.[4] There is no reason to suspect that the concept of the law of excluded middle is contained within the concept of John, for example, so while 'John is such that the law of excluded middle is true' is presumably analytic on the first conception, it is presumably synthetic on the second.

Much of The First Critique concerns the coarse-grained conception of analyticity (to the extent that the distinction is recognized), and for Kant's dialectic purposes, this was entirely appropriate. One of his primary contributions was the development of the synthetic *a priori*—claims which can be known to be true purely based on mental reflection, but whose truth does not merely depend upon the meanings of their terms (geometric truths, for example). Had Kant operated with the fine-grained conception of analyticity, this would be trivial and uninteresting. 'John is such that the law of excluded middle is true' would count as synthetic *a priori*. After all, reflection alone reveals that all objects (including John) are such that the law of excluded middle holds, yet there is no reason to suspect that the concept of the law of excluded middle is contained within the concept of John. Kant's program was thus made more ambitious (and, as a result, more important) by operating with the coarse-grained notion of analyticity.

But dialectic aims vary, and for my purposes it is the fine-grained notion which is more appropriate. Notably, Kant's two interpretations of analyticity are analogous to the two notions of entailment previously discussed. The coarse-grained interpretation of analyticity corresponds to the logical use of 'entailment,' while the fine-grained interpretations corresponds to the use of 'entailment' concerning the containment of meaning. And while the coarse-grained interpretation is often tacitly employed in discussions of epistemic closure, the fine-grained notion is better suited for these purposes.

Frege (1892) also discussed analyticity in terms of the containment of meaning. He

[3]For an argument that these interpretations are distinct, see Katz (1988). Macfarlane (2002) argues that Kant overlooked the distinction between these interpretations because of the weak logical resources available at his time.

[4]Or, perhaps, the link between subject and predicate preserves *aboutness*—see Yablo (2014). Additionally, some have argued that analyticity preserves directionality: that there is a way in which the truth of 'All bachelors are unmarried men' depends upon the meanings of 'unmarried' and 'man,' rather than 'bachelor'—see Fine (1994).

held that all denoting expressions have a *sense*, or way in which they denote. The phrase 'the morning star' denotes Venus in a way in which 'the evening star' does not—it denotes it as the object that appears in such-and-such a position in the morning, as opposed to so-and-so object that appears in the evening. Senses, Frege held, are compositional, so the meaning of 'The father of Caesar' depends, partially, upon the meaning of 'Caesar.' Additionally, he maintained that sentences are denoting terms—in particular, they denote their truth-values. So, 'Grass is green' denotes the True and 'The sky is green' denotes the False. Some sentences contain, as proper parts, other truth-evaluable expressions. For example,'It is raining and it is windy' contains both 'It is raining' and 'It is windy.' In these cases, the meaning of the compound sentences depend, at least partially, on the meanings of their truth-evaluable components. This is because the *sense* of these sentences—the ways in which they denote the True or the False—depend upon the senses of their proper parts. While many contemporary philosophers and linguists deny that sentences are denoting expressions, the compositionality of meaning remains a mainstay of contemporary philosophy of language.[5]

The relevant type of analyticity is thus a fine grained relation that obtains between a sentence and its truth-evaluable components. Although this relation is more fine-grained than classical equivalence, it is possible to investigate it systematically. Angell (1977, 1989, 2002) first provided a nonclassical logic of this type of containment, which was independently supplemented by (provably complete) semantics by Corriea (2004) and Fine (2015).[6] Although p guarantees the truth of $p \vee q$, the disjunction is not entirely relevant to the meaning of p. After all, p makes no mention of q. In contrast p is entirely relevant to the truth of $\neg\neg p$, and both p and q are entirely relevant to the conjunction $p \wedge q$. So while disjunction elimination fails on Angell's logic, double negation introduction (and elimination) and conjunction elimination are preserved. I myself am sympathetic to the details of Angell's system, but many do not bear on the current project, so I do not place inordinate weight on it. However, I assume that the axioms of double negation and conjunction elimination are compulsory for a logic of analytic containment.

This system has already been put to work. Fine (Forthcoming), for example, employs it in a theory of partial truth.[7] Some are susceptible to the idea that sentences can be partially true and partially false. Plausibly, 'Cats are mammals and dogs are reptiles' is merely partially true. Of course, merely partially true sentences are strictly false; any partially false sentence is itself false. Nevertheless, we can understand a sense in which these sentences remain partially true. It is difficult to account for merely partial truth with classical entailment. One might suggest that a sentence is merely partially true just

[5]For discussions and critiques, see Montague (1970); Horwich (1997); Szabó (2000); Fodor (2001); Johnson (2004).

[6]For the purposes of this paper, I omit the formalisms these philosophers provide. I direct those interested in the technical details to their original papers.

[7]This system has also been put to use in deontic logic (see Fine (2018*a*,*b*)) and generalized identity (see Correia and Skiles (2017); Elgin (Forthcoming)).

in case it entails something which is true and something else which is false. But, if this were so, every false sentence would count as being merely partially true—after all, every false sentence entails infinitely many trivial logical truths. One might, instead, suggest that a sentence is merely partially true just in case it entails a *nontrivial* truth and a *nontrivial* falsehood, but this excludes too much. '1 + 1 = 2 and dogs are reptiles' would not count as merely partially true, because the relevant truth (that 1 + 1 = 2) is trivial. It is much easier to account for partial truth with a fine-grained notion of analyticity. A sentence is merely partially true, on this conception, just in case it has at least one analytic part which is true and at least one analytic part which is false. This requires a notion of analyticity which comes apart from classical entailment.

As many are doubtlessly aware, some dispute the analytic/synthetic distinction. Quine (1951) canonically argues against the distinction in two ways.[8] First, he highlights the difficulty of formulating the distinction in more fundamental terms. Although someone might attempt to define analyticity in terms of synonymy or translation, these notions demand just as much clarification as analyticity does. Independently, this might appear to be a weak reason to abandon the analytic/synthetic distinction. After all, we have already encountered difficulty in formulating epistemic closure in more fundamental terms, but this is not an adequate reason to abandon closure principles. However, Quine then outlines a positive conception of language (which he expands upon in Quine (1960)) according to which the meanings of words are determined by a complex web of language. On this semantic holism, the meaning of any one term depends upon the use and meanings of all others; there is no distinctive way in which the meaning of 'bachelor' depends upon the meaning 'unmarried' yet does not depend upon the meaning of 'chocolate.' If there is no distinctive way the meanings of some terms depend upon the meanings of others, then the traditional conception of analyticity ought to be abandoned.

I doubt that Quineans will find my version of closure appealing—it crucially relies upon a distinction they reject.[9] I primarily direct this paper towards those who, like myself, hold analyticity in high regard. However, there is an avenue left open to those who reject the analytic/synthetic distinction. Quine allows for a notion of analyticity restricted to logic; there is a sense in which p analytically contains q if p logically entails q. The system Quine had in mind was classical logic, rather than Angellic logic (at the time, nonclassical systems were in their infancy), but there was no principled objection to a notion of analyticity given in purely logical terms. In theory, a Quinean could adopt a notion of analyticity restricted to Angell's logic while rejecting the connection between 'Sarah is an optometrist' and 'Sarah is an eye doctor.'

[8] For a classic response, see Grice and Strawson (1956).

[9] More precisely, Quineans ought to maintain that Analytic Closure is either ill-formed or vacuously true. If 'analytic' lacks meaning, then presumably the sentences it occurs within are also meaningless (including Analytic Closure). If, instead, it is meaningful but applies to no sentences whatsoever, then Analytic Closure is vacuously true; any agent who knows that p also knows all of the analytic parts of p, because there are no analytic parts of p.

It is my hope that the relevant notion of analyticity is, by now, sufficiently clear. In claiming that knowledge is closed under analytic content, I appeal to a fine-grained relation between sentences—a relation that obtains just in case the meaning of one contains the meaning of another. And while this notion is controversial, it remains theoretically useful. However, it is worth distinguishing my version of closure from other prominent interpretations. Some subscribe to generative versions of closure—ones that directly concern our ability to gain knowledge by recognizing consequences—like the following:

> **Generative Closure:**
> If S knows that p and S concludes that q by competently deducing it from p, then S knows that q.[10]

Generative Closure differs from Analytic Closure, and I take no stand on whether it is true (after all, I allow for the possibility that there are multiple correct versions of the closure principle). Its target differs from mine. It concerns our ability to generate knowledge: given knowledge of one proposition, how can one come to know another? In contrast, I am concerned with what antecedent knowledge one must have in order to know certain propositions. I do not claim that those who competently deduce p from a known $p \wedge q$ thereby come to know that p; rather, I claim that agents must know that p in order to count as knowing that $p \wedge q$ in the first place. So while Generative Closure might be seen a providing sufficient conditions for knowledge of some propositions, I am concerned with necessary conditions of others.

There are three reasons I believe Analytic Closure obtains: 1) it accommodates challenging cases that threaten other accounts, 2) it is theoretically useful in providing novel necessary and sufficient conditions for knowledge (while evading the preface paradox, which affects related proposals), 3) knowledge requires the recognition of meaning. Let us take these points in turn.

Conjuncts are analytic parts of their conjunctions—the meaning of $p \wedge q$ contains the meanings of both p and q. If knowledge is closed under analytic content, then everyone who knows that $p \wedge q$ also knows that p and knows that q. Similarly, the meaning of $\neg\neg p$ contains the meaning of p. If knowledge is closed under analytic content, then everyone who knows that $\neg\neg p$ also knows that p. This does not obtain because of oddities about conjunction and negation in particular, but rather because they are instances of a more general pattern of analytic containment. However, there are many classical entailments that are not instances of analytic containment. Analytic Closure does not require that everyone who knows that water is wet knows that Fermat's Last Theorem is true, despite classically entailing the theorem, because the meaning of 'water is wet' does not contain the meaning of Fermat's Last Theorem. And so, Analytic Closure accommodates negation and conjunction while avoiding the implausible result that all epistemic agents are logically

[10] For defenses of this sort of view, see, e.g., Williamson (2000).

omniscient. But conjunction and negation are the easy cases; many versions of Closure diagnose them correctly. What about the hard ones?

Suppose that John is at a zoo, and observes what he takes to be a zebra.[11] The creature satisfies every zebra-criterion he considers: it is shaped roughly like a horse, it has black and white stripes, it stands behind a sign reading 'zebra,' etc.. For these reasons, John (correctly) forms the belief 'The animal in front of me is a zebra.' It is reasonable to suppose that John knows that the animal in front of him is a zebra. Whatever the threshold for knowledge is, John has sufficient evidence and mental capacities (and whatever else may be required) to meet it. Does John know that the animal in front of him is not a painted mule? Presumably, the fact that the animal is a zebra entails that it is not a painted mule (on the classical, rather than Angellic, use of 'entailment'), so Naïve Closure entails that he does. And if John were to competently deduce that the animal is not a painted mule from the fact that it is a zebra, he would satisfy the conditions for knowledge specified by Generative Closure as well.

Intuitively, this is incorrect; John does not know that the animal is not a painted mule. After all, if a trickster had sneaked in during the night and replaced the zebra with a painted mule, John would still believe that the animal in front of him is a zebra. His belief is not appropriately responsive to truth. So it appears that the Naïve and Generative accounts misdiagnose this case. However, it poses no threat to Analytic Closure. After all, there is no reason to suspect that the meaning of 'zebra' contains the meaning of 'painted mule,' and likewise no reason to think that the meaning of 'The animal in front of me is a zebra' contains the meaning of 'The animal in front of me is not a painted mule' (or vice versa). It is readily possible to know what a zebra is without having any idea of what a painted mule is. And so, Analytic Closure does not require that John know that the animal is not a painted mule in order to count as knowing that it is a zebra.

Take another example. Suppose that Sarah is driving through the countryside and observes what she takes to be a red barn, and forms the belief 'There is a barn.'[12] As it turns out, she is correct—she does indeed observe a barn. In ordinary circumstances, many would be inclined to say that she knows that there is a barn. Suppose, however, that the field is filled with (extremely convincing) paper mâché barn facades. Sarah has been fooled into believing that these facades were barns numerous times before. In fact, there is one such facade immediately behind the barn she observes, so she would have believed that there is a barn regardless of whether or not one is there. In this case, many maintain that Sarah does not know that there is a barn—her belief is not appropriately responsive to the presence of barns.

Additionally, let us suppose that all of the barn facades are all blue, while the actual

[11] This type of example first appears in Dretske (1970). For responses, see Luper (1984); BonJour (1987); DeRose (1995).

[12] The first discussion of this sort of case I am aware of appears in Goldman (1976)—crediting Ginet. It was presented by Kripke (2011) as an argument against Nozick (1981). This case was also used to support Dretske (2003, 2005)'s account of knowledge. For a related discussion, see Hawthorne (2005).

barns are all red. In addition to believing that there is a barn, Sarah forms the belief 'There is a red barn.' Unlike her belief that there is a barn, Sarah's belief that there is a red barn *is* responsive to the presence of red barns—had she not observed a red barn, she would have formed the belief 'There is a blue barn' instead. But although her belief that there is a red barn is responsive to the presence of red barns, some maintain that Sarah does not know that there is a red barn. After all, how could it be that Sarah knows that there is a red barn, yet does not know that there is a barn, when the first so obviously entails the second?

This case is somewhat trickier than the painted mule; intuitions vary about whether Sarah counts as knowing that there is a red barn in this situation. How Analytic Closure rules depends on considerations about analyticity. If 'There is a barn' is an analytic part of 'There is a red barn,' and if Sarah does not know that that there is a barn, then Analytic Closure entails that Sarah does not know that there is a red barn either. It is a prerequisite for knowing that there is a red barn that Sarah knows there is a barn. If, however, 'There is a barn' is not an analytic part of 'There is a red barn,' then Analytic Closure has no bearing on this case.[13]

So much for difficult cases. Analytic Closure also affords the resources for novel necessary and sufficient conditions for knowledge. Care must be taken, as the distinction between proper and improper parthood bears on these conditions. Following standard discussions of mereology, we might say that improper parthood is reflexive, while proper parthood is irreflexive; no sentence is an analytic proper part of itself, but every sentence is an analytic improper part of itself. The intuitive thought is that a proper part of s must be supplemented by something else in order to yield s, while an improper part may or may not require anything else. For the conditions I have in mind, improper analytic parthood is at issue. And while I maintain that these are necessary and sufficient for knowledge, I do not take them to constitute a reductive analysis. 'Knowledge' appears on both sides of this principle, and the reason it obtains is partially, though not entirely, trivial.

> **Knowledge and Knowledge of Improper Parts (KKIP)**:
> S knows that p if and only if, for any q which is an *improper* analytic part of p, S knows that q.

Every sentence is an improper analytic part of itself. So, for a sentence p without any proper analytic parts, this condition entails that S knows that p just in case S knows that p—trivial, to be sure, but perfectly true. KKIP gains its teeth from sentences with proper analytic parts. In these cases, it entails that S knows that p if and only if S knows every analytic part of p. These parts include p, but include other sentences as well. So, for

[13] The logic of analytic containment that Angell provides is sentential, and so does not have immediate implications regarding the relation between the predicates 'is a red barn' and 'is a barn.' For my part, I am inclined to treat 'a is a red barn' as synonymous with 'a is red and a is a barn,' in which case 'a is a barn' is an analytic part of 'a is a red barn.' I leave the discussion of demonstratives like 'there' for another time.

example, if 'Roses are red and violets are blue' has three analytic parts ('Roses are red,' 'Violets are blue,' and 'Roses are red and violets are blue'), then an agent S knows that roses are red and violets are blue if and only if S knows that roses are red, that violets are blue, and that roses are red and violets are blue. If S fails to know any one of the three, then S does not know that roses are red and violets are blue.

It is straightforward to establish that Analytic Closure entails KKIP. Suppose that Analytic Closure is true, and that an arbitrary agent S knows that p (for an arbitrary p). Analytic Closure requires that S knows that all parts of p are true, so, for any q which is an improper analytic part of p, S knows that q. Suppose, instead, that an arbitrary p is such that, for any improper part q of p, S knows that q. Because improper parthood is reflexive, p is an improper analytic part of p, so S knows that p. Therefore, S knows that p if and only if, for any q which is an improper analytic part of p, S knows that q—KKIP obtains.

It is important to specify KKIP in terms of improper, rather than proper parthood not only because analytic closure does not entail the proper-parthood analogue of KKIP, but also because such a principle is clearly false. Consider the following:

> **Knowledge and Knowledge of Proper Parts (KKPP):**
> S knows that p if and only if, for any q which is a *proper* analytic part of p, S knows that q.

One direction of this biconditional is unproblematic. Analytic closure requires that if S knows that p, then for any q which is an analytic part of p, S knows that q. Problems arise for the other direction: for the claim that if S knows that q, for any q which is an analytic proper part of p, then S knows that p.

One issue concerns *the preface paradox*—a puzzle concerning the accumulation of knowledge.[14] Suppose that an author rationally believes each assertion in her book; she has checked her sources carefully, and confirmed all evidence she relies upon multiple times. As it turns out, she has made no mistakes, and, for each sentence in her book, counts as knowing that that sentence is true. Nevertheless, she has her doubts about the conjunction of all of these sentences. From her point of view, the odds that a particular sentence is wrong are minimal; the odds that there is a mistake somewhere or other are relatively high. So while she counts as knowing each sentence (we can suppose), she does not know the conjunction of all of these sentences. This poses problems for KKPP. Presumably, the only analytic parts of the conjunction of the sentences within the author's book are the individual sentences. By stipulation, the author knows that each of these sentences is true, and so, according to KKPP, the author knows that their conjunction is true as well. But this is incorrect; the author does not know this conjunction because the accumulation of doubt is sufficiently high.

[14] For the original discussion of the preface paradox, see Makinson (1965). For discussions about what agents ought to believe in light of this paradox, see, e.g., Kyburg (1961); Sorensen (2003).

In contrast, the preface paradox poses no threat to KKIP. In order for the author to know that the conjunction of the sentences in her book are true, KKIP entails that she knows that every improper part is true. This conjunction is an improper part of itself, so if she does not know this conjunction she does not know all of the improper parts of this conjunction. And so, even if the author knows that each sentence in her book is true, KKIP does not entail that she knows that their conjunction is true.

Another issue is *the problem of radical omniscience*—everyone knows absolutely everything. This problem rests on a somewhat controversial picture of language; that there are sentences without analytic parts, and that every other sentence is composed, in some way or other, out of these mereologically simple sentences. Let us suppose that this is correct, and select an arbitrary sentence s without proper analytic parts. Because s has no analytic proper parts, everyone vacuously knows that all of its proper analytic parts is true. So, KKPP entails that everyone knows that s is true. The selection of s was arbitrary, so everyone knows that every sentence without analytic proper parts is true. For example, if 'Apples are red' and 'Oranges are orange' both lack analytic proper parts, then everyone knows that apples are red and everyone knows that oranges are orange. Next, select an arbitrary sentence s' whose only analytic proper parts are sentences that, themselves, lack analytic parts. Because everyone knows that all sentences without analytic proper parts are true, everyone knows that all of the analytic proper parts of s' are true. Therefore, KKPP entails that everyone knows that s' is true. And because the selection of s' was arbitrary, everyone knows that every sentence whose only proper analytic parts themselves lack proper analytic parts are true. Returning to our previous example, if 'Apples are red' and 'Oranges are orange' both lack analytic proper parts, (and if 'Apples are red and oranges are orange' has no proper parts other than those two sentences), then everyone knows that apples are red and oranges are orange. It should be clear that this process can be continued indefinitely. As a result, KKPP entails that everyone knows that every sentence whatsoever is true. It is worth recognizing how radical this omniscience is; it is an omniscience that many theologians do not even attribute to God. KKPP does not simply entail that everyone knows that every *true* sentence is true; it entails that everyone knows that *every* sentence is true *regardless of that sentence's truth-value*. All falsehoods are known, by everyone, to be true. This result is too absurd to be worthy of further consideration; KKPP must be abandoned. By entailing KKIP rather then KKPP, Analytic Closure thus entails necessary and sufficient conditions for knowledge that are plausible, as opposed to closely related conditions that are not.

In addition, I hold that Analytic Closure poses an often-overlooked requirement on knowledge; propositional knowledge to require the recognition of meaning. In order to know that a proposition p is true, an agent must grasp what p means. In doing so, an agent will at least tacitly realize what the analytic parts of p are, and that a commitment to p is accompanied by a commitment to its analytic parts. And she understands that whatever justification and evidence they have in support of p gives justification and evidence to p's analytic parts.

I grant that there may be cases in which an agent is disposed to assert that p without understanding what p means. Suppose, for example, a layperson overhears a scientist claim 'electrons exhibit quantum entanglement' without any concept of what 'electrons' or 'quantum entanglement' mean. He might be inclined to take the scientist at her word, to report what the scientist said to his colleagues, and to even bet that electrons exhibit quantum entanglement in the right situation. Nevertheless, I deny that such a person knows that electrons exhibit entanglement—at best, he knows that 'electrons exhibit quantum entanglement' expresses something or other true. He is incapable of performing other activities often associated with knowers, like explaining what he knows in different terms, or integrating this claim with related facts about electrons or entanglement.

I close by briefly discussing the connection between Analytic Closure and epistemic contextualism: the view that claims of the form 'S knows that p' vary in truth-value depending on the context they occur in.[15] As it turns out, many contextualists are tacitly committed to the truth of Analytic Closure. Some contextualists endorse a principle like Naïve Closure within a linguistic context, but not between contexts with different standards for knowledge. For example, the inference from 'I am wearing a blue t-shirt' to 'I am wearing a t-shirt' is presumably permissible—the linguistic shift from 'blue t-shirt' to 't-shirt' does not affect the standards for knowledge. However, the assertion 'I know that I am wearing a blue t-shirt, therefore I know that I am not a brain in a vat' is false. By raising the possibility that I am envatted, I alter the context in such a way that raises the standards for knowledge.

Suppose that this is true, i.e. that Naïve Closure holds within (but not between) linguistic contexts, and further suppose that someone truly asserts 'S knows that $p \wedge q$.' In this case the assertion 'S knows that p' does not alter the linguistic context. After all, each of those terms occurred but one sentence before. And because the linguistic context is unaltered, the inference to 'S knows that p' is preserved—it is a valid inference within the same linguistic context. Just so for other instances of analytic containment. For an arbitrary context in which 'S knows that p' is true, the sentence 'S knows that q' does not change the linguistic context if q is an analytic part of p—after all, something containing the very meaning of q occurred but one sentence before. Because the context is unaltered and because the inference is valid, the inference to 'S knows that q' is universally preserved. Of course, one need not be a contextualist in order to accept a principle like Analytic Closure, but many contextualists would do well to avail themselves of it.

[15] For defenses of closure, see, e.g., Cohen (1986); DeRose (1992, 1995); Lewis (1996).

References

Angell, Richard. 1977. "Three Systems of First Degree Entailment." *Journal of Symbolic Logic* 47(1):147.

Angell, Richard. 1989. Deducibility, Entailment and Analytic Containment. In *Directions in Relevance Logic*, ed. J. Norma and R. Sylvan. Kluwer pp. 119–43.

Angell, Richard. 2002. *A-Logic*. University Press of America.

BonJour, Laurence. 1987. Nozick, Externalism, and Skepticism. In *The Possibility of Knowledge: Nozick and His Critics*, ed. Steven Luper. Rowman and Littlefield.

Cohen, Stewart. 1986. "Knowledge and Context." *The Journal of Philosophy* 83:574–83.

Correia, Fabrice and Alexander Skiles. 2017. "Ground, Essence and Identity." *Philosophy and Phenomenological Research* pp. 642–70.

Corriea, Fabrice. 2004. "Semantics for Analytic Containment." *Studia Logica* 77(1):87–104.

Davies, Martin. 1998. Externalism, Architecturalism, and Epistemic Warrant. In *Knowing our Minds*, ed. Crispin Wright, Barry Smith and Cynthia MacDonald. Oxford University Press pp. 321–61.

DeRose, Keith. 1992. "Contextualism and Knowledge Attributions." *Philosophy and Phenomenological Research* 52(4):913–29.

DeRose, Keith. 1995. "Solving the Skeptical Problem." *The Philosophical Review* 104(1):1–52.

Dretske, Fred. 1970. "Epistemic Operators." *Journal of Philosophy* 67:1007–23.

Dretske, Fred. 2003. Skepticism, What Perception Teaches. In *The Skeptics*, ed. Steven Luper. Ashgate Publishing pp. 105–18.

Dretske, Fred. 2005. Is Knowledge Closed Under Known Entailment? In *Contemporary Debates in Epistemology*, ed. Mattias Steup and Ernest Sosa. Blackwell.

Elgin, Samuel. Forthcoming. "The Semantic Foundations of Philosophical Analysis.".

Fine, Kit. 1994. "Essence and Modality." *Philosophical Perspectives* 8:1–16.

Fine, Kit. 2015. "Angellic Content." *Journal of Symbolic Logic* 45(2):199–226.

Fine, Kit. 2018*a*. "Compliance and Command I—Categorical Imperatives." *The Review of Symbolic Logic* 11(4):609–33.

Fine, Kit. 2018*b*. "Compliance and Command II—Imperatives and Deontics." *The Review of Symbolic Logic* 11(4):634–64.

Fine, Kit. Forthcoming. "Toward a Theory of Partial Truth.".

Fodor, Jerry. 2001. "Language, Thought and Compositionality." *Mind and Language* 16(1):1–15.

Frege, Gottlob. 1892. "On Sense and Reference." *Zeitschrift für Philosophie und Philosophische Kritik* 100:25–50.

Goldman, Alvin. 1976. "Discrimination and Perceptual Knowledge." *The Journal of Philosophy* 73:771–91.

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

Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford University Press.

Hawthorne, John. 2005. The Case for Closure. In *Contemporary Debates in Epistemology*, ed. M. Steup and E. Sosa. Blackwell pp. 26–43.

Horwich, Paul. 1997. "The Composition of Meaning." *The Philosophical Review* 106(4):503–33.

Johnson, Kent. 2004. "On the Systematicity of Language and Thought." *Journal of Philosophy* 101(3):111–39.

Kant, Immanuel. 1781. *The Critique of Pure Reason*. Cambridge University Press.

Katz, Jerrold. 1988. *Cogitations*. Oxford University Press.

Kripke, Saul. 2011. *Philosophical Troubles*. Oxford University Press.

Kyburg, Henry. 1961. *Probability and the Logic of Rational Belief*. Kluwer.

Lewis, David. 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* 74(4):549–67.

Luper, Steven. 1984. "The Epistemic Predicament: Knowledge, Nozickian Tracking and Skepticism." *Australasian Journal of Philosophy* 62:26–50.

Macfarlane, John. 2002. "Frege, Kant and the Logic of Logicism." *The Philosophical Review* 111(1):25–65.

Makinson, D. C. 1965. "The Paradox of the Preface." *Analysis* 25:205–7.

Montague, Richard. 1970. "Universal Grammar." *Theoria* 36(3):373–98.

Nozick, Robert. 1981. *Philosophical Explanations*. Cambridge University Press.

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

Quine, W. V. 1960. *Word and Object*. MIT Press.

Sorensen, Roy. 2003. Paradoxes of Rationality. In *The Handbook of Rationality*, ed. Al Mele. Oxford University Press pp. 257–75.

Szabó, Zoltán Gendler. 2000. *Problems of Compositionality*. Garland.

Vogel, Jonathan. 1990. Are There Counterexamples to the Closure Principle? In *Doubting: Contemporary Perspectives on Skepticism*, ed. M. Roth and G Ross. Kluwer Academic Publishers.

Vogel, Jonathan. 2000. "Reliabilism Leveled." *Journal of Philosophy* 97:602–23.

Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford University Press.

Wright, Crispin. 1985. "Facts and Certainty." *Proceedings of the British Academy* 71:429–72.

Yablo, Steve. 2014. *Aboutness*. Princeton University Press.