

Knowledge is Closed Under Analytic Content

Samuel Z. Elgin

March, 2020

Abstract

I am concerned with epistemic closure—the phenomenon in which some knowledge requires other knowledge. In particular, I defend a version of the closure principle in terms of analyticity; if an agent S knows that p is true, then S knows that all analytic parts of p are true as well. After targeting the relevant notion of analyticity, I argue that this principle accommodates intuitive cases and possesses the theoretical resources to avoid the preface paradox. I close by arguing that contextualists who maintain that knowledge attributions are closed within—but not between—linguistic contexts are tacitly committed to this principle's truth.

Knowledge of some claims requires knowledge of others. All those who know that $p \wedge q$ also know that p —as do all those who know that $\neg\neg p$. Agents ignorant of the truth of p lack the epistemic resources to know either the conjunction or double negation. In order for someone to count as knowing either logically complex expression, they must also count as knowing that p . This much is relatively uncontroversial, but there is currently no consensus on the scope of, and the foundations for, this phenomenon. Under what conditions does knowledge of one proposition require knowledge of others? Is it a feature of conjunction and negation in particular, or are they instances of a more general pattern? What is it in virtue of that this epistemic restriction occurs?

This is the interpretive question of the closure principle, the most basic formulation of which is the following:

Naïve Closure:

If an agent S knows that p , and p entails q , then S knows that q .

Naïve Closure is transparently false. I presumably know that I have hands, and the fact that I have hands may well entail that Goldbach's conjecture is true, yet I do not know whether Goldbach's conjecture is true. Agents regularly fail to recognize the logical consequences of what they know; sometimes they actively disbelieve those consequences. And, surely, if an agent disbelieves that something is true, then they do not know that it is true. We could amend Naïve Closure by requiring agent S to believe that q in order to count as knowing that q , but this does not accommodate cases in which 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 clauses, amendments and restrictions (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 be an intuitive and straightforward connection between knowledge and entailment quickly becomes an extraordinarily convoluted affair. As with other interpretive philosophical debates, it is not entirely clear that there is a unique resolution—perhaps multiple versions of closure obtain. However, concerns about an overabundance of uncontroversial interpretations may be premature, as we have yet to uncover one. Those investigating the phenomenon of closure may thus be seen as seeking *an adequate* analysis of closure, rather than *the* analysis of closure.

Despite the increasing complexity of accounts, and the arduousness of precise formulation, many continue to find some version of closure appealing.¹ Of course, there are detractors. Dretske (1970, 2005), for example, argues that knowledge involves the ability to track a proposition's truth, and 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. And, quite recently, Alspecter-Kelly (2019) argues that, regardless of what epistemic commitments we have (aside from a potential commitment to closure) the denial of closure has lucrative theoretical payoffs—shedding light on the nature of justification, epistemic bootstrapping, the transfer of warrant, and much more besides. 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 needed. What is required, instead, is a refined notion of entailment. Often philosophers (as well as mathematicians, linguists, and many others) use 'entailment' to refer to a relation which preserves truth. Proposition p entails proposition q , on this conception, just in case the truth of p guarantees the truth of q . However, there is another notion of entailment concerning the preservation of *meaning*. On this alternate conception, proposition p entails proposition q just in case the meaning of q is a part of the meaning of p . If the meaning of 'John is a bachelor' contains the meaning of 'John is male,' then the claim that John is a bachelor entails that John is male. Conversely, if the meaning of 'Grass is green' does not contain the meaning of ' $2 + 2 = 4$,' then the claim that grass is green does not entail that $2 + 2 = 4$, despite the fact that it guarantees the arithmetic truth. This, I claim, is the notion of entailment under which knowledge is closed: the type

¹For defenses of closure, see, e.g., Vogel (1990, 2000); Williamson (2000); Hawthorne (2004, 2005).

between a sentence and its analytic parts. In particular, I subscribe to the following:

Analytic Closure:

If an agent S knows that p , and q is an analytic part of p , then S knows that q .²

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 (and ordering) of its terms. Plausibly, ‘All vixens are foxes’ and ‘All doctors are doctors’ are both analytic. In contrast, a sentence is said to be synthetic just in case its truth-value depends upon more than the meanings (and ordering) of its terms. Plausibly, ‘All vixens have tails’ and ‘Some doctors are wealthy’ are both synthetic.

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 interpretation, 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,’

²To the best of my knowledge, this paper constitutes the first defense of this precise principle. However, there are precursors in the literature. Stine (1976), for example, argues that putative failures of closure (which are discussed below) occur because the conclusions are not relevant to the truth of the premises. The difference between our approaches is that, while Stine incorporates relevance into the pragmatics, my own view is that relevant implication lies at the level of semantics; we ought to understand what p means in terms of the implications it is entirely relevant to. Another account, which is probably the closest currently available in the literature, occurs in Yablo (2014, 2017). Yablo defends a version of closure in terms of *aboutness*: knowledge is closed under entailments that are about the same thing. I suspect that the decision between my and Yablo’s versions of closure ought to be determined by examining the comparative advantages of theories of subject-matter in terms of aboutness and truth-maker semantics. Any adequate such comparison would take us far afield, and warrants a paper in its own right (for one such discussion, see Rothschild (Forthcoming)). For the moment, I note one point of difference between our approaches: on Yablo’s account, what a sentence is about changes based on context. For example, if someone were to reason from ‘I have hands’ to ‘I am not a brain in a vat,’ the content of ‘I have hands’ may change in order to include the possibility of envatted brains. On my approach, in contrast, the semantic content of ‘I have hands’ is static: its meaning need not change by later raising the possibility of envatted brains. I suspect, however, that many advantages for my version of closure apply to Yablo’s account as well. The discussion remains valuable because it vindicates formulations of closure in this area (Yablo does not, for example, discuss how accounts of closure in terms of aboutness accommodate the preface paradox—nor does he argue that contextualists are committed to this principle’s truth). Additionally, Elgin (2018) appeals to, but does not defend, a principle along these lines.

³For 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.

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. 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.

To the extent that Kant recognized this distinction, he was primarily concerned with the coarse-grained notion of analyticity in *The First Critique*, and this was entirely appropriate given his dialectic purposes. 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.

However, dialectic aims vary, and it is the fine-grained interpretation is better suited for the present purpose. Notably, Kant’s two conceptions of analyticity correspond to the two notions of entailment previously discussed. The coarse-grained interpretation is analogous to the logical use of ‘entailment,’ while the fine-grained interpretation is analogous to the use of ‘entailment’ concerning the containment of meaning. And while the coarse-grained notion is often tacitly employed in discussions of closure, it is the fine-grained notion which, I maintain, is more appropriate.

Analyticity is often described as a relation between predicates (or properties). It is not uncommon, for example, for philosophers to assert that the predicate ‘male’ is contained within the predicate ‘bachelor.’ However, I am presently concerned with a relation between sentences (or propositions)—a relation that obtains between p and q just in case the meaning of q is contained within the meaning of p .⁴ There is a perfectly intelligible sense in which the meaning of ‘Roses are red and violets are blue’ contains the meaning

⁴The conception of analyticity as a relation between sentences can be traced back at least to Frege (1892), who identified the meanings of denoting terms with their *sense*—or the way in which they denote. Senses, Frege held, are compositional, so the meaning of ‘the father of Caesar’ depends, partially, on 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. And because senses are compositional, the meaning of ‘Grass is green and the sky is blue’ depends (partially) on the meaning of ‘Grass is green.’ The way in which the conjunction denotes the True depends partially upon the way that ‘Grass is green’ denotes the true. While many contemporary philosophers and linguists deny that sentences are denoting expressions, the compositionality of language remains a mainstay of contemporary philosophy of language—for discussions in this area, see Montague (1970); Horwich (1997); Szabó (2000); Fodor (2001); Johnson (2004).

of ‘Roses are red,’ and in which the meaning of ‘Jane is a vixen’ contains the meaning of ‘Jane is a fox.’ And it is worth noting, in this regard, that the sentential conception of analyticity can be straightforwardly folded into the predicative conception. Sentences can be conceived of as predicates with 0-adicity—i.e., as predicates that take no objects as their argument, so there is no obstacle to conceiving of sentential analyticity as a limiting instance of predicative analyticity.

The notion of analyticity I am concerned with is thus a fine-grained relation between a sentence and its truth-evaluable components. And while this notion of analytic entailment comes apart from classical entailment (there is, for example, no reason to suspect that the meaning of $p \vee q$ is contained within the meaning of p), it can be investigated systematically. Angell (1977, 1989, 2002) first formalized a nonclassical logic of analytic containment along these lines.⁵ On this system, the meanings of conjunctions invariably contain the meanings of their conjuncts, and the meanings of double negations invariably contain the meaning of their double negatums. So while conjunction and double negation elimination are preserved on this system, disjunction introduction fails. I myself am sympathetic to the details of this account, but many do not bear on the current project, so I do not place inordinate weight on it.⁶ However, I assume that the axioms of double negation and conjunction elimination are compulsory for a logic of analytic containment.

This system was supplemented by a semantics independently by Correia (2004) and Fine (2015), and has been put to a variety of philosophical uses. Fine (Forthcoming), for example, employs it in the service of developing a theory of partial truth. Some philosophers suspect that sentences may be (merely) partially true. Plausibly, ‘Cats are mammals and Dogs are reptiles’ is one such sentence. Of course, merely partially true sentences are strictly false; any sentence which is at least partially false counts as false overall. Nevertheless, we can understand a sense in which these sentences remain partially true.

It is challenging to construct an adequate theory of partial truth armed only with the resources of classical logic. One might suspect that a sentence is merely partially true just in case it entails something which is true and entails something else which is false. But this is too permissive: every strictly false sentence would count as being merely partially

⁵For the purposes of this paper, I omit the formalisms Angell provides. I direct those interested in the technical details to his original papers.

⁶The reason I do not place inordinate weight on this particular logic is that there is some flexibility in the underlying semantics; it is possible to construct alternate, closely related systems. As it currently stands, Angell’s system is idempotent: p is analytically equivalent to $p \wedge p$ (and $p \vee p$). However, as Fine (2017) discusses, it is possible to construct a truth-maker semantics where idempotence fails. More substantively, Angell’s original system is distributive and verifies the DeMorgan laws: $p \wedge (q \vee r)$ is analytically equivalent to $(p \wedge q) \vee (p \wedge r)$, and $\neg(p \vee q)$ is analytically equivalent to $\neg p \wedge \neg q$ (for example). It is possible, however, to construct semantics on which these inferences fail. Correia and Skiles (2017) rely upon a variant which licenses one (but not both) distribution rule, while Elgin (Forthcoming) employs a system that licenses neither. Similarly, Jago (Forthcoming) constructs a system in which the DeMorgan laws do not hold. Those who maintain that knowledge that $p \wedge (q \vee r)$ entails knowledge that $(p \wedge q) \vee (p \wedge r)$ and that knowledge that $\neg(p \vee q)$ entails knowledge that $\neg p \wedge \neg q$ may employ Angell’s original system. However, those who deny that knowledge is preserved in these cases may operate with a close variant instead.

true. After all, every false sentence entails infinitely many logical truths. One might, instead, suggest that a sentence is merely partially true just in case it entails a *nontrivial* truth and a *nontrivial* falsehood, but this sacrifices too much. The sentence ‘Grass is red and $1 + 1 = 2$ ’ would not count as being merely partially true, because the relevant true conjunct (that $1 + 1 = 2$) is trivial. It is much more straightforward to construct a theory of merely partial truth on Angell’s system. A sentence is merely partially true, on this account, just in case at least one of its analytic parts is true and at least one of its analytic parts is false. Partially because disjunction introduction fails on this type of approach, this avoids the implication that every sentence is at least partially true.⁷

As many are doubtlessly aware, some philosophers dispute the distinction between analytic and synthetic truths. Quine (1951) canonically provides two arguments against the distinction.⁸ First, he notes the difficulty in formulating the distinction precisely; although some might attempt to define analyticity in terms of synonymy or translation, these notions arguably demand just as much clarification as analyticity does. Independently, this might seem to be a poor reason to abandon the analytic/synthetic distinction. After all, we have already encountered difficulty in formulating the closure principle, but this difficulty is not itself an adequate reason to abandon the notion of epistemic closure. More broadly, this argumentative style suffers from the fact that it considers a small (and inexhaustive) portion of the potential accounts of the distinction. If someone were to provide one that Quine failed to consider, it is not obvious that his first objection would have any purchase at all. However, Quine also outlines a positive conception of meaning (which he expands upon in Quine (1960)) according to which the meanings of words are determined by a complex web of language. According to this semantic holism, the meaning of any one term depends upon the meaning of all others; there is no distinctive sense in which the meaning of ‘bachelor’ depends upon the meaning of ‘unmarried’ but not on the meaning of ‘chocolate.’ If, as Quine contends, there is no distinctive way in which 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 would find my interpretation of the closure principle particularly appealing—it crucially relies upon a distinction they reject.⁹ For this reason, I primarily direct this paper to those, like myself, who hold the notion of analyticity in high regard. However, I note that there is a path left open to the Quinean: a way to accept a notion of closure in terms of analytic content. In *Two Dogmas*, Quine allows for a notion of analyticity restricted to logic; that is, he allows for a way in which p analytically contains

⁷For other philosophical uses of Angellic content, see, e.g., Fine (2012, 2018*a,b*); Correia and Skiles (2017); Elgin (Forthcoming).

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

⁹More precisely, I suspect that Quineans would hold either that my interpretation of Closure is either ill-formed or that it is vacuously true. If ‘analytic’ lacks meaning, then presumably sentences in which it occurs (including Analytic Closure) lack meaning as well. If, however, the term ‘analytic’ is perfectly meaningful, but has an empty extension, then Analytic Closure is vacuously true. Knowledge that p entails knowledge of all analytic parts of p , because there are no analytic parts of p .

everything which p logically entails. Of course, the system of logic Quine had in mind was classical logic, rather than Angellic Content (for the perhaps justifiable reason that the system of Angellic Content was not yet formalized), but he saw no principled objection to a conception of analyticity given in purely logical terms. A Quinean could accept a version of closure restricted to Angell's logic, while rejecting the claim that knowledge of 'Sarah is an optometrist' entails knowledge that 'Sarah is an eye doctor' (and the like).

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 which holds just in case the meaning of one contains the meaning of the other. And while this relation is controversial, it remains theoretically useful. However, it is worth distinguishing the type of closure with another, perhaps more prominent, version in the literature—one which concerns our ability to generate knowledge. An example of this type of principle is the following:

Generative Closure:

If an agent S knows that p and competently deduces q from p , then S knows that q .¹⁰

Generative Closure differs from Analytic Closure, and I take no stand on whether it, or some close variant of it, is true. Its target differs from mine; while it concerns agents' ability to generate knowledge given what they already know, I am concerned with preconditions for knowledge. Rather than addressing the question 'What is S in a position to know, given that S knows that p ?' I address the question 'What conditions must obtain for S to know that p in the first place?' I do not claim that those who competently deduce p from a known $p \wedge q$ thereby come to know that p ; what I claim is that knowledge that p is required in order to count as knowing that $p \wedge q$ in the first place. Thus, while Generative Closure provides putative *sufficient* conditions for knowledge of some claims, it is my aim to provide *necessary* conditions for knowledge of others.

Why do I subscribe to Analytic Closure? Three reasons, primarily. It accommodates challenging cases that threaten other accounts, it is theoretically useful in that it possesses the resources for novel necessary and sufficient (but nonreductive) conditions for knowledge as well as the resources to avoid the preface paradox, and it arises from a plausible restriction on knowledge: that 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 both the meaning of p and the meaning of 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

¹⁰Yablo (2014) refers to this type of principle as 'transeunt closure'. For defenses of this sort of view, see, e.g., Williamson (2000). Like myself, Williamson is not concerned with closure insofar as it figures in debates over an analysis of knowledge, but rather takes it to be an independently important principle governing what it is that agents know.

$\neg\neg p$ 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 omniscient. But conjunction and negation are the easy cases—many interpretations of closure diagnose them correctly. What about the hard ones?

Consider a classic example: a woman—call her Linda—visits her local zoo, and observes what she takes to be a zebra.¹¹ The creature Linda observes satisfies every zebra-criteria she considers: it is shaped roughly like a horse, it has black and white stripes, it stands behind which reads 'zebra,' etc.. On the basis of this evidence, she forms the belief that the animal is a zebra. We may suppose that, as it turns out, she is correct. The animal she observes actually is a zebra. It is reasonable to suppose that, in this case, Linda knows that the animal is a zebra. Whatever the threshold for knowledge is, she has sufficient evidence, mental capacities, and whatever else may be required, to meet it.

Does Linda know that the creature in front of her is not a painted mule? Presumably, the fact that it is a zebra entails that it is not a painted mule (on the classical, rather than the Angelic use of 'entailment'), so Naïve Closure entails that she does. On that interpretation, knowledge that the creature is not a painted mule is a precondition for knowledge that it is a zebra. And if she were to competently deduce that it is not a painted mule from the fact that it is a zebra, Generative Closure would entail that she knows it is not a painted mule as well.

Intuitively, this is not so. Linda does not know that it is not a painted mule—nor does competent deduction from the fact that it is a zebra allow her to come to know that it is not a painted mule. After all, if someone had sneaked into the zoo the previous night and replaced the zebra with a cleverly painted mule, she would still believe that she observes a zebra. Her belief is not appropriately responsive to truth, so she fails to meet the threshold for knowledge. It appears, then, that both Naïve and Generative Closure misdiagnose this case.

However, this example poses no threat to Analytic Closure whatsoever. 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 'That animal is a zebra' contains the meaning of 'That animal is not a painted mule.' It is readily possible to know what a zebra is without having any idea of what a painted mule is. So, Analytic Closure does not require

¹¹This type of case was first discussed in Dretske (1970). For responses, see, e.g., Luper (1984); BonJour (1987); DeRose (1995).

Linda to know that it is not a painted mule in order to know that it is a zebra. Analytic Closure succeeds where rival accounts have failed.

Consider another classic example. Suppose that, following her stimulating trip to the zoo, Linda drives through the countryside. There, she observes what she takes to be a red barn and forms the belief ‘There is a barn.’¹² She is, once again, 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. Linda 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 Linda does not know that there is a barn—her belief is not appropriately responsive to the presence of barns.

Let us also suppose that all of the barn facades are all blue, while the actual barns are all red. In addition to believing that there is a barn, Linda forms the belief ‘There is a red barn.’ Unlike her belief that there is a barn, Linda’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 Linda does not know that there is a red barn. After all, how could it be that Linda 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 Linda 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 Linda does not know that there is a barn, then Analytic Closure entails that Linda does not know that there is a red barn either. It is a prerequisite for knowing that there is a red barn that Linda 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 the challenging cases. Another reason I subscribe to Analytic Closure is that it affords the resources for novel necessary and sufficient (but nonreductive) conditions for knowledge. Care must be taken when formulating these conditions, as the distinction between proper and improper parthood bears on how reasonable they are. Following standard discussions in the literature on mereology, we might say that *improper parthood* is reflexive, while *proper parthood* is irreflexive. Every sentence is an improper analytic part of itself, but no sentence is a proper analytic part of itself. The intuitive thought behind

¹²The 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).

¹³The 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.’

the notion of proper parthood is that a proper part must be supplemented by something else to yield that which it is a part of, while an improper part may require nothing else at all. The conditions I provide concern solely the notion of improper analytic parthood. 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 proper analytic parts, KKIP simply entails that *S* knows that *p* iff *S* knows that *p*—trivial, to be sure, but perfectly true. This principle gets its teeth from sentences with proper analytic parts. In these cases, it entails that if agent *S* knows that *p*, then *S* knows that all proper analytic parts of *p* are true as well. Let us suppose, for example, that ‘Roses are red and violets are blue’ has three improper analytic parts: ‘Roses are red,’ ‘Violets are blue,’ and ‘Roses are red and violets are blue.’ KKIP entails that someone knows that roses are red and violets are blue if and only if they know that each of these obtains. If someone fails to know any one of the three, then they do 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 select an arbitrary agent *S* and an arbitrary proposition *p*. Start from left-to-right: establishing that if *S* knows that *p*, then for any *q* which is an improper analytic part of *p*, *S* knows that *q*. This is an immediate consequence of Analytic Closure, which requires that knowledge is closed under analytic content. Move to right-to-left: establishing that if *S* knows that for any *q* which is an *improper* analytic part of *p*, *S* knows that *q*, then *S* knows that *p*. Recall that the notion of improper parthood is reflexive: every sentence is an improper analytic part of itself. So if *S* knows that all improper analytic parts of *p* are true, then *S* knows that *p* is true, because *p* is, itself, one of these analytic parts. Therefore, *S* knows that *p* if and only if, for any *q* which is an improper analytic part of *p*, *S* knows that *q*, and KKIP obtains.

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

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 remains innocuous. 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 aggregation 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.

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. Unlike the preface paradox, 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 general picture 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. And because everyone knows that every proper analytic part of s is true, KKPP entails that everyone knows that s is true. Of course, the selection of s was arbitrary—it could pick out any simple sentence whatsoever. Therefore, everyone knows that every sentence without proper analytic 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,

¹⁴For 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).

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—everyone is omniscient.

It is worth recognizing how radical this omniscience is; this far surpasses the problem of logical omniscience defenders of closure sometimes struggle to accommodate. This 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 than 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 requires the recognition of meaning. In order to know that a proposition p is true, an agent must grasp what it is p means.

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 (Rob) 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 Rob knows that electrons exhibit quantum entanglement.

There are activities that knowers are often able to perform, which lie beyond Rob's epistemic abilities. For example, knowers are often able to describe what it is that they know in other terms. If someone knows that their car parked on Market Street, they might convey this information by stating 'My car is parked on Market Street,' 'My Ford is one block from Island Avenue' or 'My automobile is where I left it.' This is an activity Rob is unable to engage in. Because he does not know what 'electrons exhibit quantum entanglement' means, he cannot describe it in other terms.

In addition, Rob cannot integrate the claim that electrons exhibit entanglement within his larger body of knowledge. He cannot, for example, understand how the entanglement of electrons interacts with other aspects of subatomic physics, or recognize what sorts of experiments would either confirm or disconfirm the scientist's claim. While he may be able to perform some inferences based on what he has been told, his cognitive restrictions are quite severe. It is not even clear, for example, that if the scientist were to also assert 'positrons exhibit quantum entanglement' that Rob could infer 'electrons and positrons

exhibit quantum entanglement.’ In order for Rob to justifiably make that inference, he would need to know that ‘quantum entanglement’ is not ambiguous: that the term means that same thing when applied to ‘electrons’ and to ‘positrons.’ But if Rob does not know what ‘quantum entanglement’ means, he does not know whether its meaning has shifted. For this reason, I deny that Rob knows that electrons exhibit entanglement. At best, he knows that ‘electrons exhibit quantum entanglement’ expresses something or other true.

If knowledge requires the recognition of meaning, it is straightforward to see how a principle like Analytic Closure would arise. In coming to know that p , an agent recognizes what it is that p means. In doing so, such 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. It is unsurprising, then, that knowing that p is accompanied by knowing the analytic parts of p as well.

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.

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

References

- Alspector-Kelly, Marc. 2019. *Against Knowledge Closure*. Cambridge University Press.
- 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.
- Correia, Fabrice. 2004. “Semantics for Analytic Containment.” *Studia Logica* 77(1):87–104.
- 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 Sceptics*, 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. 2018. “On Question-Begging and Analytic Content.” *Synthese* .
- Elgin, Samuel. Forthcoming. “The Semantic Foundations of Philosophical Analysis.”
- Fine, Kit. 2012. “Counterfactuals Without Possible Worlds.” *Journal of Philosophy* 109(3):221–46.
- Fine, Kit. 2015. “Angelic Content.” *Journal of Symbolic Logic* 45(2):199–226.
- Fine, Kit. 2017. Survey of Truthmaker Semantics. In *A Companion to the Philosophy of Language, Volume II*, ed. Bob Hale, Crispin Wright and Alexander Miller. Blackwell pp. 556–77.

- 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.
- Jago, Mark. Forthcoming. “Conjunctive and Disjunctive Parts.”
- 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.
- Rothschild, Daniel. Forthcoming. "Truthmakers for Semantics and Pragmatics: Presupposition and (a little) Agreement."
- Sorensen, Roy. 2003. Paradoxes of Rationality. In *The Handbook of Rationality*, ed. Al Mele. Oxford University Press pp. 257–75.
- Stine, G. C. 1976. "Skepticism, Relevant Alternatives, and Deductive Closure." *Philosophical Studies* 29(4):249–61.
- 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.
- Yablo, Steve. 2014. *Aboutness*. Princeton University Press.
- Yablo, Steve. 2017. "Open Knowledge." *Philosophical Studies* 174(4):1047–71.