

UC San Diego

UC San Diego Previously Published Works

Title

Knowledge is closed under analytic content

Permalink

<https://escholarship.org/uc/item/0kn799f2>

Journal

Synthese, 199(1-2)

ISSN

0039-7857

Author

Elgin, Samuel Z

Publication Date

2021-12-01

DOI

10.1007/s11229-021-03027-5

Peer reviewed



Knowledge is closed under analytic content

Samuel Z. Elgin¹ 

Received: 5 March 2020 / Accepted: 9 January 2021
© The Author(s) 2021

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

Keywords Analyticity · Closure principle · Entailment · Truth-maker semantics

Knowledge of some truths 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 (although not entirely) uncontroversial, but there is currently no consensus on the scope of, and on 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 .

¹ I suspect that some philosophers will deny that knowledge is closed under conjunction elimination—perhaps maintaining that it is possible for an agent to know that $p \wedge q$ while actively disbelieving that p . Such philosophers will presumably reject any version of closure analogous to the one I defend. In contrast, I maintain that cases in which an agent disbelieves that p are cases in which that agent does not really know that $p \wedge q$ (though they may be disposed to assert the conjunction). This paper is directed towards those who diagnose such cases similarly.

✉ Samuel Z. Elgin
elgin.samuel@gmail.com

¹ University of California, San Diego, San Diego, CA, USA

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. In light of these considerations, some may be tempted by the following modification:

Not-Quite-So-Naïve Closure:

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

Not-Quite-So-Naïve Closure is an improvement—but a marginal one. It faces many of the same problems that undermined Naïve Closure. Agents may fail to put two and two together: they might know that p , and know that p entails q , but yet not conclude that q on that basis. Such agents may still, it seems, disbelieve that q and so not count as knowing that q . We might further amend this principle 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 .

At this point, philosophers diverge over the appropriate response. Some abandon closure principles for knowledge entirely—maintaining that *being in a position to know*, rather than knowledge, is closed under entailment. Along these lines, Yli-Vakkuri and Hawthorne claim, “It is tempting to think that while knowledge itself does not obey any closure principles, being in a position to know does” (Yli-Vakkuri and Hawthorne, Forthcoming, pg. 1).² Others, in contrast, continue to add further clauses, amendments, and restrictions (as philosophers 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.

A particularly prominent instance of the latter approach concerns our ability to generate knowledge. An example of this type of principle is the following:

² For other examples, see Williamson (2000), Smithies (2019). It is challenging to charitably interpret Yli-Vakkuri and Hawthorne's type of claim. If, as I maintain, knowledge of conjunctions requires knowledge of conjuncts, then there must be *some* version of knowledge-closure that obtains. The question is not whether knowledge obeys any closure principles, but rather how widespread this phenomenon is. If closure were restricted to conjunction elimination, there would remain a closure principle for knowledge—albeit one with a much narrower scope than we might have expected. The most charitable readings of the Yli-Vakkuri and Hawthorne quote that I can construct are the following: either they deny that knowledge is closed under conjunction elimination (in which case it seems quite doubtful that any version of knowledge-closure succeeds), or else they hold that knowledge technically obeys closure principles, but these principles are too weak to perform substantive theoretical work.

Generative Closure:

If an agent S knows that p and competently deduces q from p , then S knows that q .³

Agents, it seems, are capable of generating knowledge by coming to recognize the consequences of what they already know. Principles like Generative Closure are intended not only to avoid the counterexamples that plagued the Naïve and Not-Quite-So-Naïve accounts, but also to reflect this capacity.

Despite the increasing complexity of accounts, and the arduousness of precise formulation, many continue to find some version of closure (whether for knowledge or for being in a position to know) appealing.⁴ 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 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, Alspectorkelly (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.

The approach I defend within this paper is comparatively conservative.⁵ I do not abandon closure outright; I do not shift to discussing closure for *being in a position to know*; I do not suggest any further clauses or amendments to existing principles. Rather, I maintain that Not-Quite-So-Naïve Closure is perfectly true as it stands. If an agent S knows that p and knows that p entails q , then S knows that q . All that is required to salvage this principle is a refined notion of entailment. Often philosophers (as well as mathematicians, linguists, and many others) use 'entailment' to refer to a relation that 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

³ Yablo (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. See, also, Becker (2017).

⁴ For defenses of closure, see, e.g., Vogel (1990), Vogel (2000), Williamson (2000), Hawthorne (2004, 2005).

⁵ There are several reasons why this approach might be described as 'conservative.' As it will become clear when this view is fully articulated, there are many cases it does not apply to. For example, (as I suggest when discussing the classic painted-mule case), it is doubtful that this version is strong enough to vindicate skeptical arguments that rely upon closure. Those who appeal to closure in defense of skepticism may take this to constitute a reason to embrace a more ambitious interpretation than I defend. In contrast, those who would resist skepticism may consider it an advantage that this version is theoretically interesting while evading the skeptical result.

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

Analytic Closure:

If an agent S knows that p and knows that q is an analytic part of p , then S knows that q .⁶

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.⁷ 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

⁶ To 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. Additionally, in Elgin (2020), I appeal to, but do not defend, a principle along these lines.

⁷ 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.

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. The law of excluded middle is not 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 that 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 the concept of the law of excluded middle is not 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 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 (classically) 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 that, I maintain, is more appropriate.⁸

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 an analytic part of the predicate ‘bachelor.’ However, I am presently concerned with a relation between sentences (or propositions)—a relation that obtains between

⁸ Chalmers (2012) offers a conception of analyticity that does not neatly align with the course-fine distinction outlined here. Building upon Carnap (1955), Chalmers argues that an intension is a function from possible cases to objects. So, for example, ‘Grass is green’ is a function from possible cases to truth-values. A sentence s is analytic for a speaker at a time just in case s has the value ‘true’ in all possible cases for that speaker at that time (Chalmers later shifts to discussions of intension in Bayesian terms). The reason this does not collapse into the coarse-grained conception is that ‘possible cases’ refers to epistemically—rather than metaphysically—possible cases. It may be metaphysically necessary that water is H_2O , but it is not epistemically necessary for all speakers that water is H_2O . After all, some speakers do not know the chemical composition of water. Chalmers does not describe a notion of analytic proper parthood—but it is reasonable to interpret him as claiming that if two sentences involve the same intension for the same speaker at the same time, then the sentences mean the same thing to that speaker at that time (he claims that there is no “conceptual change” in these cases—see pg. 203). This account remains too coarse-grained for the present purpose. It may be that $p \vee \neg p$ and $q \vee \neg q$ are associated with the same functions for a speaker at a time (if the speaker judges each to be true in every possible case)—but the meaning of p is relevant to the first disjunction, but irrelevant to the second.

p and q just in case the meaning of q is contained within the meaning of p .⁹ There is a perfectly intelligible sense in which the meaning of ‘Roses are red and violets are blue’ contains the meaning 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-arity—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.¹⁰ 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 upon it.¹¹ 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 mam-

⁹ The 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).

¹⁰ For the purposes of this paper, I omit the formalisms Angell provides. I direct those interested in the technical details to his original papers.

¹¹ The 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 that 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.

mals and dogs are reptiles' is one such sentence. Of course, merely partially true sentences are strictly false; any sentence that 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 that is true and entails something else that is false. But this is too permissive: every strictly false sentence would count as being merely partially true. After all, every false sentence entails infinitely many logical truths. One might, instead, suggest that a sentence is merely partially true just in case it entails a *nontrivial* truth and a *nontrivial* falsehood, but this sacrifices too much. The sentence 'Grass is red and $1 + 1 = 2$ ' would not count as being merely partially true, because the relevant true conjunct (that $1 + 1 = 2$) is trivial. It is much more straightforward to construct a theory of merely partial truth on Angell's system. A sentence is merely partially true, on this account, just in case at least one of its analytic parts is true and at least one of its analytic parts is false. Partially because disjunction introduction fails in this system, this avoids the implication that every sentence is at least partially true.¹²

As many are doubtlessly aware, some philosophers dispute the distinction between analytic and synthetic truths. Quine (1951) canonically provides two arguments against the distinction.¹³ 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.

¹² For other philosophical uses of Angell's Content, see, e.g., Fine (2012, 2018a, b), Correia and Skiles (2017); Elgin (Forthcoming).

¹³ For a classic response, see Grice and Strawson (1956). More recently, Wilson (2006) defends an account of meaning along similar lines to, but perhaps more radical than, Quine's—one that puts pressure on any systematic rules that govern the meanings of expressions.

I doubt that Quineans would find my interpretation of the closure principle particularly appealing—it crucially relies upon a distinction they reject.¹⁴ 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 everything that 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 that holds just in case the meaning of one contains the meaning of the other. And while this relation is controversial, it remains theoretically useful.

Why do I subscribe to Analytic Closure? Two reasons, primarily. It accommodates challenging cases that threaten other accounts, and 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. 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 $\neg\neg p$ 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?

¹⁴ More 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 .

Consider a classic example: A woman—call her Linda—visits her local zoo and observes what she takes to be a zebra.¹⁵ 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 a sign that 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.

Many people maintain that (at least when considered in isolation) this is a case in which Linda knows that the animal before her is a zebra. Whatever the threshold for knowledge is, she has sufficient evidence, mental capacities, and whatever else may be required, to meet it. But does Linda know that the animal before 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.

At this point, intuitions diverge. Some maintain that Linda does not really know that the animal is a zebra—initial appearances aside—because she does not know that it is not a painted mule.¹⁶ Such philosophers presumably form this belief because they accept some version of closure or other: one on which Linda’s ignorance that the animal is not a painted mule precludes her from knowing that it is a zebra. Others, however, retain the intuition that she knows it’s a zebra. After all, nearly everyone accepts that she knows that the animal is a zebra before the possibility of painted mules arises; we regularly ascribe knowledge to agents in more epistemically precarious positions than Linda.

What does Analytic Closure rule in this case? As always, the answer turns on considerations of analyticity. If ‘The animal is not a painted mule’ were an analytic part of ‘The animal is a zebra,’ then those who know that analytic truth must know that the animal is not a painted mule in order to count as knowing that it is a zebra. I see no reason to suspect that this is the case. 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 Linda to know that it is not a painted mule in order to know that it is a zebra.

This strikes me as a mark in favor of Analytic Closure. But, before continuing, I should note the flexibility inherent in the present account. Recall that I allow for the possibility of multiple correct interpretations of closure: it may be that multiple

¹⁵ This type of case was first discussed in Dretske (1970). For responses, see, e.g., Luper (1984), Bonjour (1987), DeRose (1995).

¹⁶ A sophisticated version of this response has been advanced by contextualists like Cohen (1986); DeRose (1992), DeRose (1995), Lewis (1996). They maintain that in ordinary linguistic contexts Linda counts as knowing that the animal in front of her is a zebra. However, in contexts where the possibility of painted mules is salient, the standards for knowledge rise. Consequently, although ‘Linda knows that the animal is a zebra’ may ordinarily be true, it is typically false when preceded by ‘Linda does not know that the animal is a painted mule.’

distinct formulations are correct. If a philosopher rules differently in this case—and maintains that Linda does *not* know that the animal is a zebra (because she does not know that it is not a painted mule), that is compatible with what I say. There may be some further closure principle that precludes Linda from counting as knowing that the animal is a zebra. In this sense, Analytic Closure does not force our hand on this debate: it is a formulation of closure that allows for the possibility that Linda knows that the animal is a zebra but does not know it is not a painted mule, but it is compatible with more aggressive formulations that do not.

When I first wrote this paper, I defended the following variant of Analytic Closure:

Naïve Analytic Closure:

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

This version, I take it, accommodates many of the cases that closure is intended to capture. It requires, for example, that knowledge is closed under conjunction and double negation elimination. It is also false. The problem is that agents may be ignorant of the meanings of terms involved in what they know. Suppose that someone hears from a doctor that their grandmother has arthritis. Such a person, plausibly, comes to know that their grandmother has arthritis on the basis of the doctor's testimony. Nevertheless, it remains possible for such a person to not know that their grandmother has inflammation of the joints on the grounds that they do not know that arthritis is inflammation of the joints. This raises problems for the previous principle—according to which “ S knows that her grandmother has arthritis” entails “ S knows that her grandmother has inflammation of the joints” on the plausible assumption that “ S 's grandmother has inflammation of the joints” is an analytic part of “ S 's grandmother has arthritis.”

This problem is closely related to a prominent view in the philosophy of language: semantic externalism. According to the externalist, the meanings of terms are determined at least partially by factors independent of speakers.¹⁷ The meaning of the term ‘arthritis,’ for example, is partially determined by the historical use of the term—rather than merely by a speaker's intension. Furthermore, it seems that agents are capable of knowing facts involving arthritis (such as the fact that their grandparent has arthritis) without knowing what arthritis is. Arguably, this could be extended from the meanings of terms to the meanings of sentences. The meaning of a sentence may (partially) be determined by factors independent of the speaker—factors that the speaker is ignorant of. And so, perhaps agents may know that p is true without knowing that an analytic part of p is true, because they do not know what that analytic part means.

Analytic Closure has the resources to avoid this problematic example. It remains possible for S to know that p without knowing that the analytic parts of p are true, in virtue of not knowing what these analytic parts are. In the present case, it remains possible for S to know that her grandmother has arthritis without knowing that her grandmother has inflammation of the joints, precisely because she does not know that arthritis is inflammation of the joints.¹⁸

¹⁷ See, for example, Putnam (1975), Davidson (1987).

¹⁸ My thanks to an anonymous reviewer for pressing me on this point.

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 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* that *S* knows to be 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 that *S* knows to have proper analytic parts. In these cases, it entails that *S* knows that these proper parts 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 (while knowing that the three are analytic parts of the conjunction), 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* that *S* knows to be 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* that *S* knows to be 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, on the assumption that *S* knows *p* to be an improper part of itself, then if *S* knows that all analytic parts of *p* are true, then *S* knows that *p* is true—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* that *S* knows to be 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 that S knows to be 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 that S knows to be an analytic proper part of p , then S knows that p .

One issue concerns *the preface paradox*—a puzzle concerning the aggregation of knowledge.¹⁹

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, then 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.

Some resist weakening closure principles due to this type of concern. In particular, Hawthorne (2004) argues that a principle in the vicinity of multi-premise closure is true.²⁰ After all, we regularly ascribe knowledge to agents who conjoin propositions that they independently know. If I know that Jack is giving a presentation at noon and I know that Jill is giving a presentation at 1:00 and, on that basis, conclude that Jack is giving a presentation at noon and that Jill is giving a presentation at 1:00, many maintain that I know that the conjunction is true. Of course, I might refrain from believing the conjunction if I take the risk of each conjunct to be too high, but an appropriately formulated version of multi-premise closure can account for that type of situation. Hawthorne may accept a principle like KKPP—aggregative concerns notwithstanding.

Another issue is *the problem of radical omniscience*—everyone knows absolutely everything. Unlike the preface paradox, this problem rests on a somewhat controversial

¹⁹ 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).

²⁰ More precisely, Hawthorne raises this point while discussing the Lottery Problem. It may be that John knows that he won't have enough money to go on vacation, and the claim that John won't have enough money to go on vacation entails that John will not win the lottery, but John does not know that he will not win the lottery. (Williamson 2000, pg. 117-8) takes a similar line while defending Intuitive Closure, according to which if S knows p_1, p_2, \dots, p_n and comes to believe that q by competently deducing it from p_1, p_2, \dots, p_n then S knows that q . (DeRose 2017, pg. 165-74), however, disagrees with both Hawthorne and Williamson on this matter.

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 are 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, 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 is 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. Indeed, 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.

This paper has consisted of an explication and defense of Analytic Closure: an interpretation of the closure principle according to which if an agent S knows that p , and knows that q is an analytic part of p , then S knows that q . This principle is conservative compared to other alternatives in the literature; its limited scope renders it more plausible than more aggressive formulations. It is my hope that the reader finds this interpretation to be both intuitive and theoretically useful when issues of closure arise.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted

by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Alsector-Kelly, M. (2019). *Against Knowledge Closure*. Cambridge: Cambridge University Press.
- Angell, R. (1977). Three systems of first degree entailment. *Journal of Symbolic Logic*, 47(1), 147.
- Angell, R. (1989). Deducibility, entailment and analytic containment. In J. Norma & R. Sylvan (Eds.), *Directions in Relevance Logic* (pp. 119–43). Berlin: Kluwer.
- Angell, R. (2002). *A-Logic*. Lanham: University Press of America.
- Becker, K. (2017). Deductive Closure Principle. *Routledge Encyclopedia of Philosophy*.
- BonJour, L. (1987). Nozick, Externalism, and Skepticism. In L. Steven (Ed.), *The Possibility of Knowledge: Nozick and His Critics*. Lanham: Rowman and Littlefield.
- Carnap, R. (1955). Meaning and Synonymy in Natural Language. *Philosophical Studiest*, 6(3), 33–47.
- Chalmers, D. (2012). *Constructing the World*. Oxford: Oxford University Press.
- Cohen, S. (1986). Knowledge and context. *The Journal of Philosophy*, 83, 574–83.
- Correia, F. (2004). Semantics for Analytic Containment. *Studia Logica*, 77(1), 87–104.
- Correia, F., & Alexander, S. (2017). Ground, Essence and Identity. *Philosophy and Phenomenological Research*, pp. 642–70.
- Davidson, D. (1987). Knowing One's Own Mind. *Proceedings and Addresses of the American Philosophical Association*, 60(3), 441–58.
- DeRose, K. (1992). Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research*, 52(4), 913–29.
- DeRose, K. (1995). Solving the Skeptical Problem. *The Philosophical Review*, 104(1), 1–52.
- DeRose, K. (2017). *The Appearance of Ignorance*. Oxford: Oxford University Press.
- Dretske, F. (1970). Epistemic Operators. *Journal of Philosophy*, 67, 1007–23.
- Dretske, F. (2005). Is knowledge closed under known entailment? In S. Mattias & S. Ernest (Eds.), *Contemporary Debates in Epistemology*. Oxford: Blackwell.
- Elgin, S. (2020). On question-begging and analytic content. *Synthese*, 197, 1149–63.
- Elgin, S. Forthcoming. The Semantic Foundations of Philosophical Analysis. *The Review of Symbolic Logic*.
- Fine, K. (2012). Counterfactuals without possible worlds. *Journal of Philosophy*, 109(3), 221–46.
- Fine, K. (2015). Angelic Content. *Journal of Symbolic Logic*, 45(2), 199–226.
- Fine, K. (2017). Survey of truthmaker semantics. In H. Bob, W. Crispin, & M. Alexander (Eds.), *A Companion to the Philosophy of Language* (Vol. II, pp. 556–77). Oxford: Blackwell.
- Fine, K. (2018a). Compliance and command i-categorical imperatives. *The Review of Symbolic Logic*, 11(4), 609–33.
- Fine, K. (2018b). Compliance and command II-imperatives and deontics. *The Review of Symbolic Logic*, 11(4), 634–64.
- Fine, K. Forthcoming. Toward a Theory of Partial Truth.
- Fodor, J. (2001). Language, thought and compositionality. *Mind and Language*, 16(1), 1–15.
- Frege, G. (1892). On sense and reference. *Zeitschrift für Philosophie und Philosophische Kritik*, 100, 25–50.
- Grice, P., & Strawson, P. F. (1956). In defense of a dogma. *The Philosophical Review*, 65(2), 141–58.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hawthorne, J. (2005). The case for closure. In M. Steup & E. Sosa (Eds.), *Contemporary Debates in Epistemology* (pp. 26–43). Oxford: Blackwell.
- Horwich, P. (1997). The composition of meaning. *The Philosophical Review*, 106(4), 503–33.
- Jago, M. Forthcoming. Conjunctive and disjunctive parts.
- Johnson, K. (2004). On the systematicity of language and thought. *Journal of Philosophy*, 101(3), 111–39.
- Kant, I. (1781). *The Critique of Pure Reason*. Cambridge: Cambridge University Press.
- Katz, J. (1988). *Cogitations*. Oxford: Oxford University Press.
- Kyburg, H. (1961). *Probability and the Logic of Rational Belief*. Berlin: Kluwer.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74(4), 549–67.
- Luper, S. (1984). The epistemic predicament: knowledge, nozickian tracking and skepticism. *Australasian Journal of Philosophy*, 62, 26–50.
- Macfarlane, J. (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, R. (1970). Universal grammar. *Theoria*, 36(3), 373–98.
- Nozick, R. (1981). *Philosophical Explanations*. Cambridge: Cambridge University Press.
- Putnam, H. (1975). The meaning of ‘meaning’. In G. Keith (Ed.), *Language, Mind and Knowledge* (pp. 131–93). Minneapolis: University of Minnesota Press.
- Quine, W. V. (1951). Two dogmas of empiricism. *The Philosophical Review*, 60(1), 20–43.
- Quine, W. V. (1960). *Word and Object*. Cambridge: MIT Press.
- Rothschild, D. Forthcoming. Truthmakers for Semantics and Pragmatics: Presupposition and (a little) Agreement.
- Smithies, D. (2019). *The Epistemic Role of Consciousness*. Oxford: Oxford University Press.
- Sorensen, R. (2003). Paradoxes of rationality. In M. Al (Ed.), *The Handbook of Rationality* (pp. 257–75). Oxford: Oxford University Press.
- Stine, G. C. (1976). Skepticism, relevant alternatives, and deductive closure. *Philosophical Studies*, 29(4), 249–61.
- Szabó, Z. G. (2000). *Problems of Compositionality*. Oxford: Garland.
- Vogel, J. (1990). Are there counterexamples to the closure principle? In M. Roth & G. Ross (Eds.), *Doubting: Contemporary Perspectives on Skepticism*. Berlin: Kluwer Academic Publishers.
- Vogel, J. (2000). Reliabilism leveled. *Journal of Philosophy*, 97, 602–23.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Wilson, M. (2006). *Wandering Significance: An Essay in Conceptual Behavior*. Oxford: Oxford University Press.
- Yablo, S. (2014). *Aboutness*. Princeton: Princeton University Press.
- Yablo, S. (2017). Open knowledge. *Philosophical Studies*, 174(4), 1047–71.
- Yli-Vakkuri, J., & Jonathan, H. Forthcoming. Being in a Position to Know. *Philosophical Studies* .

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.