Knowledge is Closed Under Analytic Content

Samuel Z. Elgin

Knowledge of some propositions requires knowledge of others. All those who know that $p \wedge q$ also know that p, as do those who know that $\neg\neg p$. Agents ignorant of the truth of p lack the epistemic resources to know either the conjunction or double negation. When does this hold? *Why* does this hold? Under what conditions does knowledge of one proposition entail knowledge of another?

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

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

Naïve closure is transparently false. After all, S might not realize that p entails q, and so not believe that q. We might amend Naïve closure with a requirement that S recognize that p entails q, but this still does not accommodate cases in which S fails to believe that q despite recognizing the entailment. We might introduce a requirement that S believes that q, but this does not account for cases in which S believes that q for faulty reasons, rather than because it is entailed by p. So, how ought we to understand the closure principle? As with other interpretive questions, it is unclear that there is a unique correct answer; perhaps multiple versions of closure are true. But concerns over an abundance of uncontroversial interpretations may be premature, as we have yet to uncover a single one. Despite the difficulty in constructing a formulation of the closure principle immune to counterexample, many continue to find something in this area appealing.[1]

I believe that Naïve Closure is very nearly true. No further conditions are needed, but a refined notion of entailment is. Rather than interpreting closure in terms of classical entailment, I maintain that knowledge is closed under *analytic entailment*—the type of entailment that holds 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.

[1] For defenses of closure, see, e.g., Vogel (1990, 2000); Williamson (2000); Hawthorne (2004, 2005). For attacks on closure, see, e.g., Dretske (1970); Nozick (1981).

Any defense of this principle requires some background discussion of the relevant notion of analyticity. Analyticity concerns the way the meanings of terms depend upon one another; it is sometimes said that a sentence is analytic just in case its truth-value can be determined from the meanings of its terms alone. These discussions originated in Kant (1781), who held that a judgment is analytic just in case its predicate concept is contained within its subject concept. Frege (1884, 1892) also understood analyticity in terms of the containment of meaning. He held that denoting terms have a *sense*—or way in which they denote. Senses are compositional, so the way in which one term denotes depends upon the way in which its denoting parts denote. He also maintained that sentences denote their truth-values, so the meanings of sentences depend upon the meanings of their constituent denoting parts. While many contemporary philosophers and linguists dispute the claim that sentences are denoting expressions, it is relatively uncontroversial that their meanings depend upon the meanings of their parts.

Some sentences contain other sentences as proper parts. The sentence 'It is raining and it is windy' contains the sentence 'It is raining,' and 'If it is snowing, then school is cancelled' contains 'It is snowing.' In these cases, the meanings of the compound sentences (at least partially) depend upon the meanings of their truth-evaluable components. Of course, this phenomenon is not peculiar to these particular sentences—the meaning of any conjunction depends upon the meaning of its conjuncts. The type of analytic parthood at issue for this paper is just this: the containment of the meaning of one sentence within the meaning of another. There may well be other varieties of analyticity—perhaps the term 'red' can be thought of as an analytic part of 'The ball is red,' but since this is not a relation between the meanings of sentences, it is not the type of analyticity at issue.

There are two notions of parthood that might be relevant to Analytic Closure. According to one (improper parthood) every sentence is an analytic part of itself. According to another (proper parthood) sentences are generally not analytic parts of themselves. This distinction can be largely disregarded here, but I note that adopting the notion of improper parthood provides the resources for moderately substantive necessary and sufficient conditions for knowledge; an agent S knows that p if and only if, for any q that is an (improper) analytic part of p, S knows that q. For sentences that lack analytic parts, this biconditional trivially obtains—it merely states that an agent S knows that p if and only if S knows that p. It gains its teeth from sentences with analytic parts, for here it requires that S knows that the analytic parts of these sentences hold as well.

Angell (1977, 1989, 2002) argued that the logic of analytic content is nonclassical. Although a sentence p entails $p \vee q$, the disjunction is irrelevant to the meaning of p. After all, p makes no mention of q. Angell formalized the logic of analytic containment, and this logic was recently afforded semantics independently by Corriea (2004) and Fine (2015). This system has already been put to philosophical work; for example, it figures crucially in Fine (Forthcoming)'s theory of partial truth. For the purposes of this paper, I will address two paradigmatic cases of analytic containment: that the meaning of a conjunction contains the meanings of its conjuncts, and that the meaning of a double negation contains the meaning of its double negatum. I direct those interested in further technical details to Angell's work. I assume that this logic adequately captures the relevant notion of analytic parthood, and assume that the distinction between analytic and synthetic truths is coherent. Those like Quine (1951), who dispute this distinction must search elsewhere for a closure principle.[2]

It is my hope that what Analytic Closure amounts to is, by this point, sufficiently clear. Knowledge that p entails knowledge of the analytic parts of p. So, for example, if 'John is male' is an analytic part of 'John is a bachelor,' then all those who know that John is a bachelor also know that John is male, and if 'That shape is a rectangle' is an analytic part of 'That shape is a square,' then all those who know that that shape is a square also know that that shape is a rectangle. However, it is worth distinguishing this view from alternate interpretations of the closure principle. In particular, some subscribe to generative versions of closure along the following lines:

> **Generative Closure:**
> If S knows that p and S concludes that q by competently deducing it from p, then S knows that q.[3]

I take no stand on Generative Closure—after all, I allow for the possibility of multiple correct closure principles. But its target is rather different from mine. It might be seen to provide *sufficient* conditions for knowledge, while I am concerned with *necessary* conditions. My claim is not that someone who concludes that p by competently deducing it from a known $p \wedge q$ thereby comes to know that p. Rather, I claim that in order to count as knowing the conjunction in the first place, she must know that the conjuncts are true.

While I endorse Analytic Closure, I deny its near-converse: the claim that if an agent S knows all of the analytic parts of p, then S knows that p. So, for example, I allow for the possibility that someone knows that p and knows that q but does not know that $p \wedge q$. My denial is primarily motivated by the preface paradox—a problem concerning the aggregation of knowledge. Suppose that Jane memorized the contents of a phone book, and thereby came to know each person's phone number. She knows Adam's phone number, Albert's phone number, etc.. Nevertheless, Jane may not know the conjunction of all phone numbers if she suspects that the odds that she made a mistake somewhere or other are too great. Each phone number is an analytic part of the conjunction of all phone numbers, so this is a case in which Jane knows all of the analytic parts of a sentence p without knowing that p. This example is perfectly consistent with Analytic Closure, which merely requires

[2]More precisely, Quineans ought to maintain that Analytic Closure either is an ill-formed expression or is vacuously true. If they take 'analytic' to be an ill-formed expression, then Analytic Closure is likewise ill-formed. If, instead, the term 'analytic' is perfectly well-formed but applies to nothing at all, then Analytic Closure is vacuously true; there are no sentences p and q such that q is an analytic part of p. Nevertheless, vacuous truth is only moderately preferable to falsity, so I suspect such philosophers will continue the search for other closure principles.

[3]See, e.g., Williamson (2000).

that if Jane *were* to know the conjunction of all phone numbers, then she would know everyone's phone number.

Why do I countenance Analytic Closure? Two reasons, primarily. The first reason is that it is extremely intuitive. Although philosophers often trade in intuitions, this virtue remains underappreciated. Many philosophical positions (even many philosophical positions described as 'intuitive') come across as moderately plausible at best, if not clearly incorrect. It is strong a mark in favor of a theory if it strikes those who consider it with intuitive force: as the sort of thing which we should have accepted all along. Strong—but inconclusive. I myself am hit by the intuitive force of naïve set theory, yet I do not believe it to be true.

Perhaps a stronger reason is that I take propositional knowledge to require the recognition of meaning—in order to know that a proposition p is true, an agent must recognize what it is that p means. In doing so, an agent will realize what the analytic parts of p are, and that a commitment to p is accompanied by a commitment to its analytic parts. And she understands that whatever justification and evidence they have in support of p gives justification and evidence to p's analytic parts.

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

I also countenance Analytic Closure because it correctly diagnoses relevant cases. Presumably, I know that I am human yet do not know whether Goldbach's conjecture is true. Nevertheless, the fact that I am human entails Goldbach's conjecture (assuming, of course, that Goldbach's Conjecture is true), so I am capable of knowing a proposition without knowing all of its entailments. This is no threat to Analytic Closure—after all, there is no reason to suspect that Goldbach's conjecture is any part of what 'I am human' means.

Or take Dretske (1970)'s often-discussed example of a painted mule. Suppose I were to go to a zoo, observe a striped mammal in front of me and, on the basis of my observation and an aptly-positioned placard, conclude 'That is a zebra.' Many suggest that I know that what I see is a zebra, but do not know that it is not a painted mule (after all, if it were a painted mule, I would still maintain that it is a zebra). So although I know that what I see is a zebra, I do not know something that it entails—that what I see is not a painted mule. This too poses no threat to Analytic Closure, so long as 'That is not a painted mule' is not an analytic part of 'That is a zebra.' I see no reason to suspect that the first sentence is an analytic part of the second—after all, it is quite possible for agents to know what a zebra is without having any conception of what mules are (something which would be impossible if the meaning of 'mule' were a constituent part of the meaning of 'zebra').

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.[4] 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 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.

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

References

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

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

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

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

Corriea, 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.

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

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

Frege, Gottlob. 1884. *Foundations of Arithmetic*. Northwestern University Press.

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

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.

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

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

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

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

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.