

Knowledge Closure and Competent Deduction

Short abstract: Knowledge closure spells out the intuition that agents can extend their knowledge by performing competent deductions from what they know. Alas, this principle leads to the puzzle of *epistemic immodesty*: If agents used deduction across the board, they could acquire knowledge that they do not seem able to acquire by those means. A traditional response is to reject knowledge closure. A less radical solution is to preserve knowledge closure but impose limits on the knowledge that can be acquired by deduction. Contextualism and contrastivism exemplify the compatibilist strategy. I shall defend a new form of compatibilism. My starting point will be the emerging consensus that *competent deduction* must figure in the antecedent of knowledge closure. I suggest that competent deduction requires the selection of premise-beliefs that are relevant to answering the question at hand, and I submit that this requirement is not fulfilled in alleged cases of epistemic immodesty.

Consider Dretske's (1970) zebra scenario:

The zebra

Hannah is in the zoo and gets a good look at a zebra in a pen clearly marked 'zebra.' Hannah has normal vision, lacks expertise in zoology, and has performed no checks on the animal. Given that the observation conditions are normal, Hannah comes to know that there is a zebra in the pen.

Suppose now that Hannah deduces that there is an animal before her, that there is something with black and white stripes in the pen, and so on. On the plausible assumption that deduction is a good way of extending knowledge, Hannah comes to know all these propositions. This intuition is captured by 'knowledge closure':

Knowledge Closure

For all propositions, p , q , and agents S : If S knows p , and S competently deduces q from p (thereby coming to believe q while retaining her knowledge of p), then S knows q .

Suppose that Hannah were to make deductions across the board. She could reason that, since there is a zebra in the pen, there is not a cleverly disguised mule in the pen. Deductions of this sort give rise to the puzzle of *epistemic immodesty*: Knowledge closure seems to imply that Hannah can acquire knowledge that she does not seem able to acquire.¹

Several philosophers have sketched analyses of knowledge (attributions) that impose restrictions on the knowledge that can be acquired by deduction. Thus, they preserve knowledge closure and avoid epistemic immodesty.² These solutions are ‘compatibilist’. I shall explore a new form of compatibilism. My starting point is this: the logical consequences of what we know are cases of knowledge *only if* they are reached by *competent deduction*. I will argue that the competent-deduction requirement is not fulfilled in the alleged cases of epistemic immodesty.³ If the agent cannot competently deduce the logical consequences of what she knows, knowledge closure is vacuously satisfied in the problematic cases.

What does it take for deduction to be competent? Virtue epistemology provides a fruitful framework to address this question. Imagine an archer’s shooting of an arrow. Suppose that the speed and orientation of the arrow, as it leaves the bow, would take it to the bull’s-eye in normal conditions. In that case, the shot is *adroit*. Another possibility is that the shot is *accurate* but incompetent. Suppose that, at the time of release, the arrow tip points in a direction that would not take the arrow to the bull’s-eye in normal conditions. Still, when the arrow leaves the bow, a powerful magnet attracts it to the bull’s-eye. In that case, the shot is accurate but incompetent. Finally, a shot can be both adroit and accurate but not *apt*. Imagine that a gust of wind takes the arrow off course, and then another gust puts it back on course. If the shot ends up hitting the bull’s-eye, it will be accurate, but it won’t be accurate through the exercise of the archer’s competence (Sosa 2017: 72-3). There is an explanatory relation between competence and accuracy. When a performance is competent, it *tends* to lead to an accurate outcome in normal conditions. As Sosa (2017: 191) puts it: “A competence is a disposition (ability) to succeed when one tries.” This explanatory relation is also captured by aptness. When the performance is apt, the causal responsibility for its accuracy is creditable—to a significant extent—to a disposition seated in the agent.

¹ See Unger (1975: 24-5) and Cohen (2002: 313-4).

² This includes contrastivism (e.g., Schaffer 2005, 2007) and most forms of attributer contextualism (e.g., Cohen 1988, 1999; DeRose 1995; Lewis 1996).

³ Hereafter, I abbreviate ‘the alleged cases of epistemic immodesty’ to ‘the problematic cases.’

A similar explanatory relation holds for competent deduction. If a deduction is competent, it tends to achieve its goal in normal conditions. Moreover, the causal responsibility for achieving its goal is creditable—to a significant extent—to a disposition seated in the agent. Our theory of competent deduction will depend on what we take the aim of deduction to be. Plausibly enough, deduction has a conditional aim: Given true-premise beliefs, deduction endeavors to provide true conclusion-beliefs. Let us call this ‘the alethic aim.’ Still, agents typically try to achieve the alethic aim in the context of inquiry. In inquiry, deduction serves the aim of answering a question. Let us call this ‘the inquisitive aim.’ The inquisitive aim is compatible with the alethic aim. After all, striving to answer a question is a good way of acquiring true beliefs.

The alethic aim asks for a specific type of competence: a disposition that is causally responsible for truth-preservation in the transition from premise-beliefs to conclusion-beliefs. In the case of deduction, a plausible candidate is the ability to apply logical rules. This yields a thin conception of deductive competence:

Thin deductive competence. A deduction, *D*, is *thinly competent* if and only if the agent manifests the disposition responsible for the application of logical rules.

If a competent logician applies logical rules to known propositions, she will be disposed to succeed relative to the alethic aim of engaging in truth-preserving transitions. But we have seen that deduction is also governed by the inquisitive aim of answering a question. If our agent merely applied logical rules to known propositions, she would not be disposed to succeed relative to the inquisitive aim. Consider the following thought experiment.

Reasoning from randomly selected premise-beliefs

A group of scientists has made a computer backup of Hannah’s propositional knowledge. They have encoded all her knowledge in short English sentences. This backup is updated whenever Hannah learns something new, forgets something or changes her mind on a given topic. In compensation for her cooperation, the scientists have decided to give Hannah access to her stored knowledge.

Unfortunately, Hannah’s knowledge is so vast that the scientists have not managed to design an efficient retrieval mechanism. Thus, although Hannah has access to all her knowledge, her access to that knowledge is random: there is a button she can press that randomly selects a piece of knowledge, displaying it on a screen. After retrieving a couple of funny known propositions, Hannah decides to play a game. She asks Peter to ask her questions. She then presses the button to randomly

retrieve a known proposition and uses that proposition to derive an answer to the question. To make the game more interesting, Hannah decides to resist any temptation to reason by relying on any 'intuitive' considerations. Instead, she strives to derive her answers by only applying logical rules to the retrieved known proposition.

The thin view of deductive competence requires only a disposition to apply logical rules to the contents of our attitudes. The thought experiment illustrates what it would mean for reasoning to manifest *only* that competence. A disposition to apply logical rules would not dispose Hannah to answer questions, even if she started from *known* propositions. Two considerations justify this assessment. First, Hannah's body of knowledge is huge. Thus, the likelihood that she will retrieve a known proposition that is relevant to answering the question at hand is very low. On Sosa's view, "[a] competence is a disposition (ability) to succeed when one tries." Having a competence to apply logical rules to one's known propositions is not sufficient to have a disposition to successfully use those known propositions to answer the question at hand. Second, even if Hannah did in fact retrieve a known proposition that had the potential to answer the question at hand, it is unclear whether only by manipulating the proposition in accordance with logical rules she could answer her question. After all, Hannah's sole reliance on logic might lead her to reason in circles or in the wrong direction.⁴

In sum, achieving the inquisitive aim of answering a question seems to be an unlikely outcome for someone who can only apply logical rules to known propositions. If reasoning seeks to answer a question, we need a thick conception of competent deduction. On the thick view, competent deduction requires a disposition to select premise-beliefs by their capacity to answer questions.

Thick deductive competence. A deduction, *D*, is *thickly competent* if and only if the agent manifests the disposition responsible for the application of logical rules and the disposition responsible for selecting premise-beliefs by their capacity to answer questions.

The thick conception predicts that the deductive-competence requirement is not satisfied in the problematic cases. If the thick conception is true, the problematic cases are not

⁴ These arguments are reminiscent of the frame problem in cognitive science. For two influential discussions, see Dennett (1998) and Fodor (2000).

counterexamples to knowledge closure. The competence argument supports this conclusion:

The competence argument

Premise 1. If there is competent deduction in the problematic cases, all the dispositions underlying competent deduction are manifested in the problematic cases.

Premise 2. Not all the dispositions underlying competent deduction are manifested in the problematic cases.

Conclusion. So, there is not competent deduction in the problematic cases.

Premise 1 is a consequence of our characterization of competent deduction as underwritten by different dispositions (e.g., to apply logical rules and to select premise-beliefs by their capacity to answer questions). Premise 2 says that not all the dispositions underlying competent deduction are manifested in the problematic cases. My claim is that the disposition responsible for premise-belief selection is not manifested in the problematic cases. Consider an epistemically immodest argument:

The mule argument

Premise 1. There is a zebra in the pen [from perception].

Premise 2. If there is a zebra in the pen, there is not a cleverly disguised mule in the pen.

Conclusion. So, there is not a cleverly disguised mule in the pen.

Consider the question of whether there is a cleverly disguised mule in the pen. A competent reasoner would not answer that question by reasoning through the mule argument. A *perceptually based* premise-belief is not a good way of answering a question about the instantiation of a property that outstrips the limits of perception. Notice that the mule argument bears a structural similarity to the thought experiment that motivated the concept of thick deductive competence: its premises seem to be randomly selected and not by their capacity to answer the question of whether there is a cleverly disguised mule in the pen. Finally, this diagnosis explains our uneasiness with the mule argument. We find ourselves hesitating between the undeniable fact that the argument is valid and the obvious fact that we cannot use that argument to answer the question at hand.

In my presentation, I shall motivate the inquisitive aim of reasoning, I shall show that the competence argument is consistent with some widespread hypotheses from cognitive science, and I will respond to some potential objections.

References

- Broome, J. (2013) *Rationality Through Reasoning*, Malden (MA.), Wiley-Blackwell
- Cohen, S. (1988) How to be a Fallibilist, *Philosophical Perspectives* 2: 91-123
- Cohen, S. (1999) Contextualism, Skepticism, and the Structure of Reasons, *Philosophical Perspectives* 13: 57-89
- Cohen, S. (2002) Basic Knowledge and the Problem of Easy Knowledge, *Philosophy and Phenomenological Research* 65: 309-29
- Dennett, D. C. (1998) Cognitive Wheels: The Frame Problem in AI, in his *Brainchildren: Essays on Designing Minds*, Cambridge (MA.), MIT Press: 181-205
- DeRose, K. (1995) Solving the Skeptical Problem, *Philosophical Review* 104: 1-52
- Dretske, F. (1970) Epistemic Operators, *Journal of Philosophy* 67: 1007-23
- Fodor, J. A. (2000) *The Mind Doesn't Work That Way*, Cambridge (MA.), MIT Press
- Harman, G. (1986) *Change in View*, Cambridge (MA.), MIT Press
- Lewis, D. (1996) Elusive Knowledge, *Australasian Journal of Philosophy* 74: 549-67
- McHugh, C. and J. Way (forthcoming) What is Reasoning? *Mind*, DOI: 10.1093/mind/fzw068
- Schaffer, J. (2005) Contrastive Knowledge, in: T. S. Gendler and J. Hawthorne, eds., *Oxford Studies in Epistemology* 1, Oxford, Oxford University Press: 235-71
- Schaffer, J. (2007) Closure, Contrast, and Answer, *Philosophical Studies* 133: 233-55
- Sosa, E. (2011) *Knowing Full Well*, Princeton / Oxford, Princeton University Press
- Sosa, E. (2017) *Epistemology*, Princeton / Oxford, Princeton University Press