

Knowledge Closure Without Epistemic Immodesty

Santiago Echeverri

Swiss National Science Foundation

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 deductive means. A traditional response is to reject knowledge closure. A less radical strategy is to preserve knowledge closure but impose limits on the knowledge that can be acquired by competent deduction. Contextualism and contrastivism exemplify this compatibilist strategy. I introduce a new form of compatibilism that does not incur the theoretical costs of either contextualism or contrastivism. My starting point is the emerging consensus that *competent deduction* must figure in the antecedent of knowledge closure. Next, I introduce an account of competent deduction. On this view, competent deduction requires the selection of premise-beliefs that are relevant to answering the question at hand. This account predicts that the alleged cases of epistemic immodesty violate the competent-deduction requirement.

1

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.

Hannah has normal logical skills. Hence, she can perform deductions from what she knows: she can deduce 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 these propositions. This intuitive idea 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 by deductive means.¹

Several philosophers have sketched analyses of knowledge that impose restrictions on the knowledge that can be acquired by deduction. Thus, they preserve closure and avoid epistemic immodesty.² These solutions are therefore 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.

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

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.’

2

Let us distinguish reasoning from argument (Broome 2013; Harman 1986; McHugh and Way 2018). Reasoning is an epistemic activity; an argument is an abstract object. Whereas reasoning involves attitudes such as beliefs, arguments deal with contents. In a valid argument, a premise p entails a conclusion q . This may give one a reason to believe q . Nevertheless, this is not always the case. If one has independent reasons to reject the conclusion, it may be reasonable to suspend judgment on that conclusion or disbelieve the premise. Following Broome (2013), we can term ‘premise-beliefs’ the beliefs one reasons from and ‘conclusion-beliefs’ the beliefs one reasons to.

Deduction is a type of reasoning. What does it take for deductive reasoning to be competent? The answer to that question depends on what we take the aim of deductive reasoning to be. Virtue epistemology provides a fruitful framework to flesh out this suggestion.

Imagine an archer’s shooting of an arrow. According to Sosa (2011, 2017), the archer’s performance can be assessed along three dimensions:

Adroitness: The shot is *adroit* if and only if it manifests competence.

Accuracy: The shot is *accurate* if and only if it hits the target.

Aptness: The shot is *apt* if and only if it is accurate because competent.

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. Imagine now that a gust prevents the arrow from hitting the bull’s-eye. In that case, the shot is adroit but inaccurate. Another possibility is that the shot is accurate but incompetent. Imagine 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; still, it won't be accurate through the exercise of the archer's competence (e.g., 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. This explanatory relation holds for competent deduction. If a deduction is competent, it tends to achieve its goal in normal conditions. If it is competent, the causal responsibility for achieving its goal is creditable—to a significant extent—to a disposition seated in the agent.

What is the aim of deductive reasoning? Deductive reasoning seems to have a conditional aim: Given true-premise beliefs, deductive reasoning endeavors to provide true conclusion-beliefs. Let us call this 'the alethic aim.' Agents typically try to achieve the alethic aim in the context of inquiry. In inquiry, deductive reasoning serves the aim of answering a question. Let us call this 'the inquisitive aim.' The inquisitive aim is not incompatible with the alethic aim. After all, striving to answer a question is a way of acquiring true beliefs.

Given true-premise beliefs, reasoning endeavors to provide true conclusion-beliefs. This 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 deductive reasoning, 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 the contents of her attitudes, she will be disposed to succeed relative to the aim of engaging in truth-preserving transitions. But we have seen that reasoning is also governed by the inquisitive aim of answering a question. If an agent merely applied logical rules to the contents of her attitudes, 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 backup of Hannah's propositional knowledge. They have encoded her knowledge in English sentences. In compensation for her cooperation, the scientists have decided to give Hannah access to her stored knowledge. Alas, 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 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 formulate questions. She then presses the button to retrieve a known proposition, and uses that proposition to derive an answer to Peter's question. To make the game more interesting, Hannah resists any temptation to rely on any 'intuitive' considerations. She strives to derive her answers by strictly applying logical rules to the retrieved proposition.

Our thought experiment illustrates what it would mean for deductive reasoning to manifest only *thin* deductive competence. Having only a disposition to apply logical rules to known

propositions would not dispose Hannah to answer questions. Two considerations justify this assessment. First, Hannah's body of knowledge is huge. Thus, the likelihood that she retrieves a known proposition that is relevant to answer a question is extremely 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 does not dispose one to successfully use the retrieved propositions to answer questions. Second, even if Hannah did in fact retrieve a proposition that had the potential to answer the question at hand, it is implausible to hold that she could attain her aim of answering the question *only by manipulating the known proposition with logical rules*. After all, Hannah's exclusive reliance on logic might lead her to reason in circles or in the wrong direction.

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 the contents of her attitudes. 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 that are relevant to answering the question at hand.

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 that are relevant to answering the question at hand.

3

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 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 deductive reasoning are manifested in the problematic cases.

Premise 2. Not all the dispositions underlying deductive reasoning 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 select premise-beliefs by their capacity to answer questions). Premise 2 says that not all the dispositions underlying 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.

Ordinary agents *never* reason through the mule argument to answer a question. One might think that ordinary agents lack imagination. But this overlooks another possibility, articulated by the competence argument. If a perceptually based belief does not easily come to our minds when we consider the question whether there is a cleverly disguised mule in the pen, that is *prima facie* evidence that we lack the disposition to use that premise-belief to answer the question. This verdict is supported by three considerations.

First, someone who wanted to know whether there is a cleverly disguised mule in the pen would not satisfy her curiosity by reasoning through the mule argument. Intuitively, if a belief in premise 1 is solely based on perception, it cannot answer a question about the instantiation of properties that outstrip the limits of perception. Second, 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. Third, 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 the answer the question at hand.

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
- DeRose, K. (1995) Solving the Skeptical Problem, *Philosophical Review* 104: 1-52
- Dretske, F. (1970) Epistemic Operators, *Journal of Philosophy* 67: 1007-23
- 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 (2018) What is Reasoning? *Mind* 127 (505): 167-96
- 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