

Cogency and Context

Cesare Cozzo¹ 

© Springer Science+Business Media Dordrecht 2017

Abstract The problem I address is: how are cogent inferences possible? In § 1 I distinguish three senses in which we say that one is “compelled” by an inference: automatic, seductive-rhetorical and epistemic compulsion. Cogency (in my sense) is epistemic compulsion: a cogent inference compels us to accept its conclusion, if we accept its premises and we aim at truth. In §§ 2–3 I argue that cogency is intelligible if we consider an inference as a compound linguistic act in which several component acts (assertions and hypotheses) are related to one another by a commitment that the premises support the conclusion. Non-automatic inferences are primarily public acts in an intersubjective context. But cogency arises in special contexts described in § 4, epistemic contexts, where the participants care for truth, i.e. are intellectually virtuous. An inference is cogent in an epistemic context if it stands up to all the objections raised in the context. In § 5 I consider three different kinds of cogent inferences. In § 6 I argue that in all three cases cogency is fallible and propose a fallibilist variety of inferentialism. In § 7 I distinguish context of utterance and contexts of evaluation. Cogency is relative to epistemic contexts of evaluation. However, validity, i.e. stable cogency, is transcontextual.

Keywords Inference · Cogency · Context · Inferentialism · Fallibilism

1 Three Kinds of Inferential Compulsion

Cogency is the force, or strength, of arguments. A cogent inference is an inference that “compels us to accept the conclusion” if we accept the premises. An ambiguity lurks here. In what sense are we compelled? There are at least three senses. The *first* kind of compelling inferences can be illustrated by an example (see Kahneman 2012, p. 19). We have some given data: the picture of a woman’s face. The data are automatically processed: the expression on her face is such that we immediately think, or rather feel, that the woman is angry. The given data, the picture of a face, are the premise of an inference and the processed data are the conclusion: the thought that the woman is angry. The term “inference” is now used to denote a transition from given data to processed data. But the process is so rapid that we are not even aware of a transition taking place. We see that there is a woman’s face before us, and at the same time we feel and believe that the woman is angry. In this sense we can say that the inference is unconscious, involuntary and automatic. We *cannot help* but feel that the woman is angry. The passage from premise to conclusion is not something that we *do*. It is not an act performed by an agent, but an event that simply happens to us. An event that we undergo, beyond our rational control. Precisely because the transition is automatic we are *compelled* to feel that the woman is angry. This is the first kind of compulsion: *automatic compulsion*. As Daniel Kahneman says, such automatic processes may depend on “innate skills” or may “become fast and automatic through prolonged practice” (p. 22). Automatic inferences are made continuously when we are awake; their role in our mental life is extremely pervasive and influential. However, they often lead to errors. For example, from “A bat and a ball cost \$1.10 in total” and “The bat costs \$1 more than the ball” we infer “The

✉ Cesare Cozzo
cesare.cozzo@uniroma1.it

¹ Department of Philosophy, Sapienza University of Rome, Rome, Italy

ball costs 10 cents” (see Frederick 2005). Such errors can be corrected. But in order to correct the errors arising from automatic inferences we must keep them under rational control, which involves a different kind of inferences.

The term “inference” is also used to denote acts that are not automatic: they are experienced as deliberate acts in which we consciously and voluntarily pass from the premises to the conclusion. Some of these inferences are compelling: an agent who initially does not accept the conclusion is somehow forced to change his/her mind. But these non-automatic inferences can be compelling in two ways, the second and the third kind of compelling inferences. The *second* kind of inferential compulsion can be termed *seductive-rhetorical compulsion*. An ancient description of this kind of compulsion is given by the rhetorician and sophist Gorgias. In his *Encomium of Helen* Gorgias tells that the Trojan prince Paris made a speech in order to persuade Helen to be unfaithful to her husband and sail to Troy (cf. Gorgias 2005). The speech was so compelling that Helen consented to her own abduction. When Helen reached the conclusion that she would sail to Troy, her consent was a deliberate act. Clearly, this was not an automatic inference. But Gorgias argues that Helen is not to be blamed and does not deserve her bad reputation precisely because Paris’ speech was so forcefully persuasive. Helen was, in Gorgias’ sense, *compelled* to accept the conclusion and to act accordingly: she was unfaithful because she succumbed to the *power* of the speech. The power of Paris’ speech was based on emotions. We can imagine: the fascination exerted by the beautiful words, and by the handsome speaker, the joy awakened by the expression of his love, the illusory hopes of a happy life in Troy raised by his praise and promises. Such emotions motivated her choice. Gorgias explicitly says that the compelling force of the speech has nothing to do with truth. The characteristics of seductive-rhetorical compulsion can be thus summarized: acceptance of the conclusion of the inference is not automatic, it is a voluntary and conscious act, and the determining motivations are certain emotions and desires (love, hate, pity, admiration, pride, fear, anger, greed, etc.) that do not have any prevailing relation to truth-seeking.

According to the unfavourable picture of the sophists proposed by Plato, sophists did not care at all for truth. But Socrates and Plato did care. In Plato’s dialogues, seductive rhetorical compulsion is often contrasted with a different kind of compulsion that is essentially related to knowledge and truth, an *epistemic compulsion*, which I here call *cogency*. One of Plato’s dialogues is the *Gorgias* (in Plato 1997, pp. 791–869), where the-character-Socrates confronts the-character-Gorgias, who has come to Athens in search of students to whom he can teach rhetoric. They discuss the nature of rhetoric. Socrates highlights the difference between the speeches of an

orator and the arguments of arithmetic and geometry. The persuasion instilled by the speeches of an orator provides conviction without knowledge. By contrast, the compelling arguments in arithmetic or geometry give us persuasion and knowledge. This is the *third* kind of compulsion, *cogency*. Unlike seductive compulsion, cogency fundamentally depends on truth-seeking and being rational. Aiming at truth is aiming at knowledge. Cogency is *epistemic compulsion*. An epistemically compelling inference is a transition from premises to conclusion that we must accept if we aim at truth and intend to be rational in our pursuit. Epistemic compulsion is not limited to arithmetic and geometry. A general conception of rationality originates in the *Gorgias* and in other Platonic texts, according to which obedience to cogent inferences is constitutive of rational behaviour. Socrates is the archetype of this rational attitude. In Plato’s *Crito* (46b), Socrates says “I am the kind of man who listens to nothing within me but the argument that on reflection seems best to me”, and he later adds that the argument that seems best to us is an argument that *constrains* us. The idea of a fundamental connection between rationality and obedience to cogent inferences is ancient, but still very much alive. In a recent paper Dag Prawitz (2005, p. 677) describes the compelling force of deductively valid inferences in these words: “on pain of irrationality, one must accept the truth of [the conclusion], having accepted the truth of the [premises]”.

The distinction between seductive-rhetorical compulsion and epistemic cogency is a central theme in Plato’s dialogues. Three features of cogency are thrown into relief: (1) there is a choice: one can listen to the cogent argument, but one can also decide not to listen; (2) the cogent argument is an argument that “seems best”; (3) one must accept the cogent argument, if one aims at truth. The first feature makes it clear that the special compulsion of cogent inference presupposes a special freedom. In a sense we are compelled, in a sense we are not compelled. It is not impossible to disobey reason. As we often say: “some people do not listen to logic”. This special freedom was perhaps what Wittgenstein meant in a controversial passage of *Remarks on the Foundations of Mathematics* (I, 113): “But am I not compelled, then, to go the way I do in a chain of inferences?”—Compelled? After all I can presumably go as I choose!” (Wittgenstein 1956, p. 33). David Pears explains this passage thus: “it is up to me whether I choose to follow an ordinary chain of inferences in the customary way, whereas it is not up to the locomotive to choose to stay on the rails [...] The compulsion that I feel when I follow a chain of inferences in the normal way is self imposed [...] It is not like the compulsion that keeps a locomotive on the rails” (Pears 2008, p. 87). A first approximation to an analysis of cogency is that an inference is cogent for me if, and only if, it is such that I must accept it if I aim at truth. It is

up to me whether I choose to aim at truth, but if I make this choice, then I am forced to accept a cogent inference.

The crucial choice concerns the kind of relationship that we intend to establish with an interlocutor. Cogency manifests itself in certain intersubjective contexts and not in others. Socrates makes this clear when he says to Gorgias that their discussion can go on only if they are both the kind of men “who would be pleased to be refuted” (*Gorgias* 458a) because they seek truth. The context in which cogent inferences are at home is a context of truth-seekers, where the participants are interested in the truth about a subject under investigation. The seductive arguments of an orator have their place in other contexts: a law court, a council meeting or an assembly, or verbal exchanges where the participants “are eager to win instead of investigating the subject under discussion”. Thus it seems that cogency is intelligible only if we consider inferences in an intersubjective context. This is the idea that I wish to develop in the next section.

2 Intersubjective Contexts and Categorical Assertions

Non-automatic inferences are primarily public acts in an intersubjective context and arguments are acts of concatenation of such inferences. Contexts are constituted by agent-speakers who perform linguistic acts. Speakers can play two roles: proponent and evaluator. A particular case is when proponent and evaluator coincide, but they are usually different agents. The proponent performs a linguistic act, thereby making proposals to the evaluators and undertaking commitments. The evaluator evaluates the proposal that is made. The result of the evaluation is that the evaluator accepts the linguistic act or not. The proponent undertakes commitments towards the evaluators. Commitments can be honoured or not. If it is shown that a commitment is not honoured, the evaluator has the right to criticize the proponent’s act and to refuse acceptance. These are general features of linguistic acts. They take a more specific form for different specific kinds of linguistic acts. I will now consider four kinds of acts: categorical assertions, hypotheses, hypothetical assertions, and inferences.

Luigi makes the *categorical assertion*: “that wall is three meters high”. What does Luigi do beyond uttering the sentence “that wall is three meters high”? He implicitly undertakes the commitment that the token sentence be true in the context of utterance and proposes that the interlocutor accept it as true. (It is necessary to specify “token sentence” and “in the context of utterance” because the sentence involves indexical ingredients, whose reference depends on the context of utterance: the present tense and the demonstrative “that wall”). In general we can say that a speaker *P* makes a categorical assertion of *S* if, and only if,

(1) *P* utters the sentence *S*, (2) *P* implicitly undertakes the commitment that *S* be true in the context of utterance and (3) *P* proposes to the evaluator *E* that he or she accept *S* as true.

A reasonable question is: what is it for the proponent to undertake *the commitment* that the token sentence *S* be true? We can answer by indicating certain intersubjective practices, certain normative practices that display an acknowledgment of rights and duties. The proponent undertakes the commitment that the asserted sentence be true if, and only if, two conditions are fulfilled. The first condition is that the proponent acknowledges an evaluator’s *right* to criticize the act of assertion and to demand evidence to support it, in other words the right to ask for a *justification* of the assertion. The second condition is that the proponent acknowledges his/her own *duty* to retract the assertion if the assertion is not justified and he/she is not capable of replying to the evaluator by providing sufficient evidence.

Another reasonable question is: what is it for an evaluator to *accept* a categorical assertion? The answer describes the norms for acceptance and the pragmatic consequences of acceptance. An evaluator *E* ought to accept the assertion if it is adequately justified. An evaluator *E* accepts the categorical assertion of *S* made by *P*, if, and only if, as a reaction to *P*’s utterance, *E* treats the token of *S* in the context of utterance as a reliable basis for voluntary actions (including other assertions). Suppose that Maria accepts Luigi’s assertion. What will she do? She may make inferences, drawing conclusions from the sentence that she accepted as true. But she will also perform non-linguistic actions. We can imagine that she wants to build a bookcase for that wall. She will make a certain number of shelves on the basis of the accepted sentence “that wall is three meters high”. If it later turns out that the wall is not three meters high and the number of shelves is wrong, she will have the right to blame Luigi. In general, the evaluator *E* acquires a right to attribute part of the responsibility for actions performed on the basis of the assertion to the speaker who made the assertion.

It is essential to assertoric practice that we are aware that speakers act on the accepted assertions (cf. Dummett 1981, pp. 302, 355). Under basic assertoric circumstances the hearer can take advantage of the assertoric act, benefiting from a piece of information that she, or he, would not otherwise have acquired. The hearer can thus orient her, or his, actions by relying upon the accepted assertion. On the other hand, these basic features of assertion, and the tacit agreement that an asserter is committed to the truth of the asserted sentence, are the background against which speakers can deceive others and deliberately contravene the constitutive truth-commitment. An asserter’s aim is often only to influence the hearer’s actions. A liar or a person who

simply does not care for truth can use assertion in order to make the hearer do and believe what appears to be expedient. An asserter in this case may be unwilling to retract an assertion made, even if the interlocutors raise plausible objections. The asserter can try to hide and play down any evidence against the assertion and to persuade other speakers to accept the assertion by all means, quite regardless of truth. There are many intersubjective contexts in which asserters are not concerned for truth. But even the most shameless liar realizes that if his, or her, lack of scruples is too evident, the aim of the lie will not be achieved. A liar must be capable of providing some presumed justification to support the misleading assertion, because even a liar undertakes a truth-commitment, though he, or she, does not honour the commitment. These situations are intelligible, and possible, only against the background of a basic agreement among speakers that an asserter is committed to the truth of the asserted sentence and that an evaluator ought to accept an assertion that is adequately justified. Without such a background the hearer would have no reason to rely upon the asserter's utterance, and the utterance would have no influence on the hearer's actions.

3 Inferences and Hypotheses in Context

A simple kind of inference is the following. Luigi makes a first categorical assertion "the wall is three metres high". Then Luigi makes a second assertion "The space between shelves is 23 cm". Then a third assertion: "The thickness of each shelf is 2 cm". So, Luigi has made three categorical assertions. These three acts are the premises of the inference. After performing these three acts Luigi passes to a fourth categorical assertion, the conclusion: "There are 12 shelves". This is a fourth act. But our description of the inference is not complete. An inference is not a mere temporal succession where one first asserts some premises and then a conclusion. In passing from the three premises to the conclusion Luigi takes a responsibility: he undertakes the commitment that the premises rationally support the conclusion. This commitment can be expressed by words or phrases like "therefore", "hence", "because". What we have described is a basic inference that can be understood without resorting to the notion of hypothesis. But there are other inferences whose premises are hypotheses. Thus we must explain what a hypothesis is.

Let us imagine that Luigi makes a *hypothesis*. Luigi says: "Suppose that the wall is three meters high". He is not undertaking any commitment that a sentence be true. He is simply proposing to investigate the commitments that in the context of utterance would follow from a commitment to the truth of the sentence "the wall is three meters high". In what way would they follow? Through inferences. In

general a proponent P makes the hypothesis H if, and only if, P utters H and thereby proposes to investigate the further commitments to which a categorical assertion of H would lead by acceptable inferences (in the context of utterance). These further commitments would *depend* on the hypothesis H .

When Luigi makes the hypothesis "the wall is three meters high" he does not commit himself to the truth of this sentence, because it is only a hypothesis. Starting from the hypothesis he can then perform an inference similar to that described above, and reach the conclusion "There are 12 shelves". In this case a truth-commitment concerning the conclusion is derived. But it depends on a hypothesis. Thus Luigi will not make a categorical assertion of the conclusion. Luigi's assertion of the conclusion will be hypothetical. Considering that there may be more than one hypothesis, we can explain hypothetical assertion in general as follows. A proponent P performs an act of *hypothetical assertion* of a sentence S depending on a set of hypotheses $\{H^1, \dots, H^n\}$ if, and only if, in uttering S , P proposes to accept a commitment to the truth of S (in the context of utterance) as a commitment that would derive (by acceptable inferences) from commitments to the truth of the hypotheses H^1, \dots, H^n .

Now we can remark that the premises and the conclusion of an inference need not be categorical assertions. They can be hypothetical assertions. In general a proponent P makes an *inference* if, and only if, (1) P makes (categorical or hypothetical) assertions, or hypotheses (and these acts are the premises of the inference), (2) P passes from the premises to a (categorical or hypothetical) assertion (which is the conclusion), (3) P implicitly undertakes the commitment that the premises rationally support, and thus justify, the conclusion. This explains what kind of act an inference is. An inference is a compound act in which several component acts (assertions and hypotheses) are related to one another. The relation is established by the act of undertaking a commitment that the premises support the conclusion. A reasonable question is: what is it to undertake the commitment that the premises support the conclusion? An answer can be given by indicating certain normative practices that display an intersubjective acknowledgement of rights and duties. To undertake the commitment that the premises support the conclusion is (1) to acknowledge the evaluator's right to criticize the act of passing from the premises to the conclusion (and possibly to demand a further argument to justify the transition) and (2) to acknowledge the proponent's duty to answer the criticism or otherwise to retract the inference.

The next question is: what is it for the evaluator to accept the inference? The answer describes the norms for acceptance and the pragmatic consequences of acceptance. An evaluator ought to accept an inference that is adequately

justified, and if the evaluator accepts an inference and its premises, the evaluator accepts the conclusion. But the premises can be hypothetical assertions or hypotheses. Thus acceptance of the premises can depend on a set Γ of n hypotheses $\{H^1, \dots, H^n\}$. The conclusion will depend on k hypotheses in Γ . The k hypotheses on which the conclusion depends can be less than n , because in passing from the premises to the conclusion the proponent can get rid of some hypotheses. In other words: some hypotheses can be *discharged*, so that the conclusion does not depend on them. This happens, for example, when we infer by implication introduction “if A, then B” from a premise B depending on the hypothesis A. Our conclusion “if A, then B” (in symbols “ $A \rightarrow B$ ”) no longer depends on the hypothesis A, which has been discharged.

Discharging hypotheses is an additional act that can be performed in making an inference. Another additional act is binding variables. The typical inference in which variables are bound is universal generalization. Universal generalization is a kind of inference that was already well known to the ancient Greeks and has been much discussed upon in the course of the history of philosophy (see Cellucci 2009). The proponent who performs an inference of universal generalization starts from an argument Δ^1 to the effect that a generic individual of some kind has the property expressed by F. The statement that the generic individual has F is the premise. The conclusion is that every individual of the relevant kind has the property F. In Gentzen’s natural deduction style of formalization, universal generalization is \forall -introduction (introduction of the universal quantifier \forall). The individual *free* variable “ a ” represents the indeterminate generic individual, about which no particular (undischarged) hypothesis has been made in Δ^1 , so that in Δ^1 no specific information has been provided that might distinguish a from other individuals of the same kind. The conclusion of Δ^1 , and premise of the inference of \forall -introduction, is $F(a)$ (a has property F). The conclusion of the inference is $\forall xF(x)$ (for every x , x has F), where $F(x)$ is obtained from $F(a)$ by substituting the *bound* variable “ x ” for every occurrence of “ a ”. This is the binding operation: by the inference of \forall -introduction a new argument Δ^2 has been constructed, whose immediate sub-argument is Δ^1 ; in the new argument Δ^2 the indeterminacy that characterized the conclusion $F(a)$ of Δ^1 is eliminated, because the conclusion $\forall xF(x)$ of Δ^2 does not concern an indeterminate individual, but is a determinate assertion concerning all individuals of the relevant kind. The variable “ a ”, which was free in $F(a)$, has become a *bound* variable “ x ” in $\forall xF(x)$.

What determines an *individual inference*? An inference is individuated by an *agent* who performs the act at a certain time and place (i.e. in a context of utterance) and by *seven components* (some of which may be empty). The

conclusion **C** can be supported by some non-linguistic evidence **NL** (for example acts of perception or a drawing), and by linguistic premises in a sequence **PR**, backed by some arguments in the corresponding sequence **AR**. Variables **VAR** can be bound by the inference, and hypotheses in a list **H** can be discharged. The seventh component is the commitment **S** that the conclusion is supported in a conclusive or defeasible way. These individuating ingredients are represented in a septuple $\langle \mathbf{C}, \mathbf{NL}, \mathbf{PR}, \mathbf{AR}, \mathbf{VAR}, \mathbf{H}, \mathbf{S} \rangle$, where only **C** and **S** are never absent (Cf. the notion of argumentation step in Cozzo 1994, pp. 59–63, which is a generalization of Prawitz’s notion of inference in Prawitz 1973, p. 228). In general all seven components must be mentioned: one can have the same premises and conclusion in two inferences that are different with respect to one of the other five components so that in the same context one of the two inferences is cogent and the other is not.

4 Cogent Inferences

What is a cogent inference? From Plato’s dialogues we have learned that cogency arises in special contexts, intersubjective contexts where the participants care for truth, or aim at truth. An inference may have the characteristics that could make it cogent, it can be an instance of a principle that is considered absolutely indubitable, but an interlocutor who does not care for truth is free to ignore it and will not be compelled by that inference. So, we must ask: what is it to care for truth? What is it to aim at truth?

We must answer this question in order to understand what cogency is, what it is to justify an assertion, or to undertake the commitment that an asserted sentence be true. The notion of a justification and the notion of a cogent inference are connected, because assertions are justified by means of concatenations of cogent inferences. Both notions are connected to the notion of truth-commitment, because justifications are acts aimed at supporting truth-commitments. These three connected notions, truth-commitment, justification and cogent inference, are *normative* notions: the proponent who, by making an assertion, undertakes a truth-commitment must acknowledge the duty to provide a justification of the assertion if the evaluator demands it; if an assertion is justified, or an inference is cogent, the evaluator has the duty to accept it. This connection thus reveals another feature that characterizes cogency. Unlike automatic compulsion and seductive-rhetorical compulsion, cogency is a kind of *normative* compulsion. If an inference is cogent, an evaluator ought to accept it, because an evaluator who has to decide whether to accept an assertion ought to aim at truth.

What is it to aim at truth? The best answer, I believe, consists in describing certain ways of acting. To aim at

truth is to act in accordance with intellectual virtues. In the *Gorgias*, Plato highlights an intellectual virtue when he insists on the willingness to take objections seriously and to retract one's assertions, if the objections against them are compelling (cf. *Gorgias* 458a). The first explicit formulation of the general concept of intellectual virtue is presumably offered by Aristotle in the *Nicomachean Ethics* (Aristotle 2002, p. 178, 1139b). Today virtue epistemology is an influential area of philosophical research (cf. Axtell 2000; Fairweather and Zagzebski 2001; Roberts and Wood 2007; Baehr 2011). This is not the place to examine the various conceptions of intellectual virtues, but in brief we can say that intellectual virtues are ways of acting whose conscious motivation is the search for knowledge of relevant truth. Curiosity, reflectiveness, attentiveness, carefulness, sensitivity to detail, a critical attitude, self-scrutiny, intellectual honesty, fairness, flexibility, open-mindedness, intellectual non-conformism, tenacity and intellectual courage are some of the intellectual virtues. All the intellectual virtues have a common aim, knowledge of relevant truth, but they are very diverse ways of acting. They may come into conflict with one another. Open-mindedness may clash with a critical attitude, carefulness may clash with courage. A special wisdom (another intellectual virtue) may therefore be necessary to find a way to harmonize intellectual virtues in a particular epistemic context. This is possible because a commonality unites intellectual virtues. They all "radiate" from a central core: a concern for truth (Code 1984, p. 34).

If we accept the notion of intellectual virtue, we can define "epistemic context". An epistemic context is individuated by four ingredients: (1) A community (or simply a group) of intellectually virtuous persons (proponents and evaluators of linguistic acts); (2) shared languages; (3) accepted knowledge-claims (arguments, assertions) and non-linguistic evidence; (4) problems (questions considered to be relevant and hypotheses for their solution). Epistemic contexts are subject to two kinds of pressure. The first is a negative pressure. One of the primary intellectual virtues is being critical. A critical attitude leads one to criticize old arguments and languages and to set aside old problems. On the other hand, a positive pressure is exerted by the intellectual virtues of curiosity, open-mindedness, flexibility and courage which, together with the old problems that have remained open, lead the community to new problems, new arguments, new assertions and new languages. An obvious consequence is that epistemic contexts are necessarily dynamic and unstable. An epistemic context C^1 soon generates a new epistemic context C^2 . Context C^2 will generate a further context C^3 , and so on. A *research path* is a series of epistemic contexts $C^1, C^2, \dots, C^n, \dots$, where each context C^{k+1} is generated by the preceding context C^k .

In each epistemic context, inferences are proposed and evaluated. Some of them are cogent. When is an inference

cogent? Under what conditions will the intellectually virtuous participants in context C be epistemically compelled to accept an inference? They will be compelled when their common inquiry forces them to accept the inference. A participant proposes the inference and undertakes the commitment that the premises support the conclusion. But do they really support the conclusion? The truth-seeking evaluators may criticize the act of passing from the premises to the conclusion and may demand a further argument to justify the transition. In an epistemic context, the proponent must answer the criticism, perhaps by providing a more detailed argument. If in the end the inference stands up to all criticisms and objections, the evaluators, being intellectually virtuous, will be rationally forced to accept it. A plausible analysis of cogency is thus the following: an inference I is *cogent* in an epistemic context C , if, and only if, in context C inference I stands up to all the objections raised in C . This is cogency in general. But there are at least three specific ways in which an inference can be cogent in an epistemic context.

5 Three Kinds of Cogent Inferences

The *first kind of cogent inferences* are those inferences that are treated as strongly obvious in an epistemic context. By "strongly obvious" inference I mean an inference with the following properties: the participants in the epistemic context are inclined to consider a reasonable doubt or challenge regarding the inference impossible; in spite of this inclination on the part of the community, the fact remains that doubt or challenge may arise, but if a doubt or challenge does present itself, the participants tend to explain this deviant behaviour as caused by a lack or failure of understanding on the part of the objector, or perhaps as resulting from some extrinsic factor that interfered with the exercise of the objector's cognitive faculties or intellectual virtues (for example, if the objector has taken a drug). Thus the participants neither acknowledge the possibility, nor the need of giving any justification of the inference by means of an argument consisting of more elementary steps.

Why are some inferences treated as strongly obvious? Many philosophers have answered by attributing a special kind of knowledge to the reasoners: intuitive knowledge. For example Descartes (1985, pp. 14–5) says that we see the validity of elementary inferences by mental intuition [*intuitus mentis*]: "the conception of a clear and attentive mind, which is so easy and distinct that there can be no room for doubt about what we are understanding". Descartes belongs to a philosophical tradition according to which intuitions are mental states, or acts, of immediate, non-empirical knowledge. Exponents of this tradition maintain that strongly obvious inferences (the first kind

of cogent inferences) are immediately known to be valid through intuitions.

However, the notion of intuition is controversial. What can intuitions be? How are they possible? One might suggest that inferences based on intuitions are automatic inferences in the sense of § 1. Intuitions would therefore be psychological events underlying automatic inferences. We would thus have a psychological analysis of intuition, which would explain one kind of cogency in terms of automatic inference. It is indeed true that strongly obvious inferences can sometimes be almost automatic because they are instances of rules to which a prolonged practice of rule-following is associated, and an action can become rapid and automatic from prolonged practice. This psychological explanation of strongly obvious inferences, however, faces two objections. In the first place, strongly obvious inferences are conscious and voluntary. They are never completely automatic. Second, in the epistemic context strongly obvious inferences are taken to be conducive to the attainment of truth, unlike merely automatic inferences, which as we know often lead us to make mistakes. An intuition underlying the inference concerning the bat and the ball in § 1 is not a piece of knowledge. The problem remains: how can intuition, however we understand it, make an inference *cogent*?

A more plausible answer to the question concerning the nature of strongly obvious inferences is that strongly obvious inferences are meaning-constitutive. If the participants in the epistemic context could borrow Michael Dummett's words, they would say that a strongly obvious inference "is compelling just because, presented with it, we cannot resist the passage from premises to conclusion without being unfaithful to the meanings we have already given to the expressions employed in it" (Dummett 1978, p. 301). The basic idea is that the words of a language are always linked by inferential connections and an understanding of a word always includes a grasp that acceptance of that fragment of language involves acceptance of some inferences: the strongly obvious inferences associated with that word. That is why objections to a strongly obvious inference are discarded: they are normally considered evidence that the objectors do not understand some words. Suppose that we accept English as our language and understand the word "uncle". Suppose that an interlocutor calls into question an inference from the premise "J. is Ned's uncle" to the conclusion "J. is male". We would suspect that our interlocutor does not know what "uncle" means in English. The inference is constitutive of the meaning of "uncle", therefore an English speaker who understands and accepts the word "uncle" must accept the inference. A different example, which concerns a logical connective, is implication-introduction (see above § 3). If we understand implication in the language of a non-relevant logic (classical logic or

intuitionistic logic) we know that all instances of implication-introduction are meaning-constitutive: we can infer " $A \rightarrow B$ " from a premise B, and if B depends on a hypothesis A, we can discharge hypothesis A. Observe, however, that A may be irrelevant. If B does not depend on A, we can still infer " $A \rightarrow B$ " from B without discharging any hypothesis. This inference too is meaning-constitutive and if we accept classical or intuitionistic logic, we will treat it as strongly obvious.

But an inference can be cogent in a context *C* even though it is not treated as strongly obvious. In other words, the inference does stand up to all the objections raised in the epistemic context *C*, but if it somehow happens that an interlocutor refuses to accept the inference and objects to it, the proponent's reaction will not be to treat such a reluctance to accept the inference as showing that the interlocutor does not understand or does not accept some word and the presupposed fragment of language. There are two ways in which an inference can satisfy this description.

An inference can be accepted without objection in an epistemic context, and thus be cogent, even though it is not meaning-constitutive. If someone refused to accept the inference and objected to it, the proponent would try to persuade the reluctant interlocutor by providing a more articulated argument of some kind. As a matter of fact, however, the proponent does not spell out the argument, because he, or she, and the whole epistemic community, take it for granted that though the argument can in principle be given, it is not necessary to do so. Inferences of this kind are treated as obvious in a weak sense: it is simply the case that no one calls them into question. We can term them "merely uncontroversial" inferences. William Thurston's remark concerning mathematical practice was that "within any field, there are certain theorems and certain techniques that are generally known and generally accepted. When you write a paper, you refer to them without proof" (Thurston 1998, p. 346). One might object that this simply shows that the mathematicians' behaviour is not intellectually virtuous. The objection is wrong. Excessive scrupulousness is not an intellectual virtue, it is a vice. If one always demanded fully articulated proofs, the result would be epistemic paralysis. Most real-life arguments cannot be fully articulated. If they were, they would be extremely long and complicated and therefore unintelligible. If you seek truth, you must avoid epistemic paralysis, therefore you must avoid the vice of excessive scrupulousness. Intellectual wisdom requires that the participants in an epistemic context distinguish merely uncontroversial inferences from delicate inferences. The investigators in the context may think (albeit fallibly) that with merely uncontroversial inferences the risk of error is slight. With delicate inferences the risk of error is high. For merely uncontroversial inferences detailed argument is considered neither necessary, nor useful. Merely

uncontroversial inferences are *the second kind of cogent inferences* in an epistemic context.

Delicate inferences, on the other hand, require careful treatment. For delicate inferences the virtue of intellectual caution demands criticisms, which must be answered by detailed arguments. Inferences that are delicate in an epistemic context C are controversial in C . They are challenged. Objections are raised against them. But if they stand up to all the objections advanced in C , they are cogent in C . So this is *the third kind of cogent inferences*, which we can call “successful controversial” inferences. An example is the inference from the premise “ x is a map drawn on paper” to the conclusion “ x can be coloured with only four colours in such a way that countries sharing a common border have different colours”: it became a successful controversial inference in mathematical epistemic contexts after 1977, when Kenneth Appel, Wolfgang Haken and John Koch used a computer to prove the four-colours-theorem.

6 Fallible Cogency

The dynamic nature of epistemic contexts clashes with the idea that an inference enjoying the property of cogency will never lose it. Cogency is fallible. In this section I advocate a fallibilist conception of cogency. In Socrates’ words, cogent inferences are inferences that “seem best” in the epistemic context. The participants in an epistemic context check whether inferences are good and lead the investigators nearer to truth. But the result of their critical investigation is provisional. Cogent inferences are often *unstable*. When an epistemic context C_1 generates a new epistemic context C_2 , inferences that are cogent in C_1 may cease to be cogent in C_2 . It is clear how this can happen when a delicate inference that is controversial but successful in C_1 does not stand up to a new objection raised in C_2 . So cogent inferences of the third kind can easily cease to be cogent.

Now consider an inference from the premises “Every animal is a body” and “Every man is an animal” to the conclusion “Some body is a man”. This is a syllogism in *baralipton*, which was a cogent inference for a logician in the Middle Ages. In a medieval epistemic context nobody would have challenged the inference. However, it would not have been treated as a strongly obvious inference. We can imagine that in the medieval context a logician, Petrus, proposes the syllogism in *baralipton*, and a participant in the epistemic context, let’s say Iohannes, unusually enough, calls the validity of this syllogism into question. If this had happened, an intellectually virtuous Petrus would not have reacted by simply saying that Iohannes did not understand some word and therefore his challenge should be immediately discarded without a detailed reply. Rather, the logician Petrus would

have tried to explain why the syllogism preserved truth. He would have done so by using other principles on the basis of a hidden assumption that the extensions of the terms (“Animal”, “Body”, “Man”) are not empty. So an explanation, a more detailed argument, would have been deemed possible, though normally superfluous. In the medieval context, the syllogism in *baralipton* is a “merely uncontroversial” inference and belongs to the second kind of cogent inferences. Today it is *not* cogent for us, if we follow classical (or intuitionistic) first order logic. Therefore at a certain point in the research path leading from the medieval context to our epistemic context, the syllogism in *baralipton* ceases to be cogent. Cogent inferences of the second kind can cease to be cogent.

Some readers may be inclined to think that inferences belonging to the first kind of cogency, strongly obvious inferences, are not fallible: they are cogent and remain cogent because they are meaning-constitutive. If so, the reader is inclined to adopt justificatory inferentialism. *Inferentialism* is the view that some inferences are meaning-constitutive. *Justificatory inferentialism* adds to this idea the claim that meaning-constitutive inferences are analytically valid, valid by virtue of meaning: if one understands the words whose meaning is shaped by those inferences, one must see that the inferences are valid (cf. Boghossian 1997). The meaning-constitutive inferences create a meaning that makes them valid. In this sense they are self-justifying. Nothing other than understanding is needed to make them cogent. Therefore they cannot cease to be cogent. This is justificatory inferentialism. Many philosophers, who favour or oppose inferentialism, believe that an inferentialist must be a justificatory inferentialist. It is not difficult to state the grounds for this conviction. The essential characteristic of a meaning-constitutive inference is that (if no extrinsic factor interferes with cognitive faculties or intellectual virtues) objections to the inference are discarded as simply showing that the objector does not understand some word. The immediate reaction is: “but this is what the word means!” In conclusion: the meaning-constitutive inference must be valid and preserve truth by virtue of the very meaning that it shapes. Is this conclusion really correct, or have we perhaps neglected an important possibility?

The assumption behind such a conclusion is that the activity of meaning-shaping is immune to criticism: the members of a community may wrongly believe that they have given meaning to an expression when, in reality, they have failed to do so. However, if they did succeed in shaping a meaning, the meaning that they have shaped cannot be wrong in any epistemic sense and cannot be rationally criticized; the members of the community are free to choose the meanings of their words. If this were true, meaning-constitutive inferences would be beyond rational critique. But is it true?

Our problem is whether the social institution of a meaningful language can be rationally criticized. Are there cases in which reasonable objections to meaning-constitutive inferences do not result from mere misunderstandings? Is it possible for an intellectually virtuous evaluator to criticize a meaning-constitutive inference even though he/she fully understands the words involved? The answer given by the history of science is in the affirmative. Eighteenth century zoologists knew perfectly well that in their classificatory language an understanding of the word “mammal” involved endorsing an inference from “ x is a mammal” to “ x is viviparous”. Nevertheless, they eventually accepted the reasonable objections levelled against this concept of mammal after the discovery of the female-platypus. The outcome was a new zoological classification and a new meaning for the word “mammal”. The history of zoology and botany abounds with classificatory languages that did not stand up to criticisms and were rationally replaced by other classificatory languages. Someone who understands a word can rationally criticize inferences that are constitutive of the meaning of that word. In such instances, what the critic criticizes is *the very meaning* shaped by those inferences. The critic criticizes *the fragment of language* to which the word is linked through those meaning-constitutive inferences. (An attempt at precisifying the notion of fragment of language is in Cozzo 1994, Ch. 3, § 17). In a given epistemic context, a fragment of language can be criticized for various reasons. It can be criticized because it clashes with recalcitrant empirical evidence, like the classificatory language into which it was so difficult to fit the platypus. It can be criticized because it is paradoxical, like the language of the Newtonian method of fluxions, or the language of Cantorian set theory. It can be criticized because, like the phlogiston-language, when compared with the new chemical language proposed by Lavoisier, it appears to be a less epistemically fruitful and less precise way of imposing order on experience. The activity of meaning-shaping is a risky and fallible enterprise: meanings can be rationally criticized and can turn out to be wrong. The social institution of language is no less subject to rational critique than other social institutions. Some inferences are constitutive of the meanings of words belonging to a fragment of language and in an epistemic context the language-fragment is successfully criticized. The result is that in the new epistemic context the language-fragment is abandoned and the corresponding meaning-constitutive inferences cease to be cogent. Thus cogent inferences of the first kind are fallible and can cease to be cogent.

A *logical* rule which is meaning-constitutive can also be rationally criticized. We have seen that standard implication-introduction is constitutive of the meaning of non-relevant implication. In a context C_1 where classical logic is accepted, an instance of non-relevant

implication-introduction is strongly obvious, and thus cogent. A relevant logician understands the language of classical non-relevant logic very well, but thinks that the classical way of codifying deductive reasoning is inadequate. The acceptance of logical principles depends on an epistemic cost-benefit analysis which involves various, sometimes conflicting, criteria like simplicity, elegance, fruitfulness, intelligibility, agreement with pre-existing practice (see Cozzo 1998 and Williamson 2017). Suppose that the relevant logician’s criticisms of non-relevant implication are successful (for an idea see Read 1988; Mares 2004) and suppose that those criticisms lead to the abandonment of the language of classical non-relevant logic in a new epistemic context C_2 . In context C_2 an instance of non-relevant implication-introduction, though still understood, will cease to be cogent. Thus a logical inference can be meaning-constitutive without being cogent. Cogent logical inferences of the first kind are also fallible and can cease to be cogent in a new epistemic context.

My conclusion is that justificatory inferentialism is wrong. I propose *fallibilist inferentialism*. Fallibilist inferentialism can be summarized by three theses: (1) some inferences are meaning-constitutive; (2) meaning-constitutive inferences can be criticized and rationally rejected; (3) a criticism levelled at an inference constitutive of the meaning of a word W calls into question the rational acceptability of the language-fragment which W presupposes.

7 Context of Utterance, Contexts of Evaluation, Validity

Two notions of context are relevant to our analysis of cogency. An individual inference is an act performed by an agent at a certain time and place, which is *the context of utterance*. For each individual inference there is *one* context of utterance, which is one of the factors individuating that inference. But the same individual inference can be evaluated in many different epistemic contexts. There are *many contexts of evaluation* in which evaluators can refer to individual inferences performed in the context to which they belong or in other contexts, and can discuss, criticize, accept and reject those inferences. The same individual inference may be cogent in one context and fail to be cogent in another context. In 1704 in his *Nouveaux essais sur l’entendement humain* (Livre IV, Ch. VII, § 10) Leibniz inferred “ $2+2=4$ ” from the definitions of the natural numbers 2, 3 and 4 (Leibniz 1990, p. 326). This individual inference was cogent in Leibniz’s context of evaluation. But in *Grundlagen* (Frege 1884, § 6), Gottlob Frege examined Leibniz’s inference and discovered a gap: so the inference was not cogent in the epistemic context of evaluation to which Frege belonged.

Thus cogency is relative to the context of evaluation. The participants in different epistemic contexts C_A and C_B can evaluate the same inference I made in a context of utterance C_U . But inference I can be cogent relative to C_A and not cogent relative to C_B . Suppose that the participants in the two different contexts of evaluation speak English and attach the notion of cogency expounded in this essay to the word “cogent”. Then it would be the case that the assertoric utterances of the sentence “Inference I is cogent” made in one context of evaluation C_A would be true and the assertoric utterances of the same sentence made in the other context C_B would be false. A token of the sentence would be true uttered in context C_A and another token of the same sentence would not be true in context C_B , though both tokens would refer to the same act of inference I performed in C_U . In other words, the sentence “Inference I is cogent” has the characteristic feature of indexical sentences: the truth value of its tokens changes when the sentence is uttered in different contexts, where the inference in question is differently evaluated. Statements attributing cogency to inferences behave as indexical statements.

Should we conclude that the only normative notion that can distinguish good and bad inferences is cogency relative to a context of evaluation? This question concerning the act of inference is strictly connected to a corresponding question concerning the act of assertion: should we admit that the only normative notion that can distinguish good and bad assertions is the notion of justifiedness? The two questions are connected because an inference (in the basic case) is an act of passing from certain categorical assertion-premises to a categorical assertion-conclusion. In performing an inference, the speaker’s aim is to provide rational support for the conclusion, i.e. a justification of the conclusion. By drawing the conclusion, the proponent makes an assertion, and thus undertakes a commitment that the asserted sentence is true in the context of utterance. We have seen that the nature of the truth-commitment manifests itself in the proponent’s acknowledgement of an evaluator’s right to ask for a justification of the assertion (cfr. § 2). So, it seems that we have two normative notions concerning an assertion: truth of the asserted token sentence in the context of utterance and justifiedness of the assertion. Are these equivalent notions? Of course, one might simply say that they are not equivalent because truth is applied to a token sentence and justifiedness is applied to an act of assertion. But even if what is liable to be justified or unjustified is the act of assertion, and not the asserted token sentence, we can define a notion of assertibility that is applicable to token sentences and in a sense corresponds to the justifiedness of an assertion: a token sentence S asserted in a context of utterance C_U is assertible if, and only if, the assertion made by uttering S is justified. Are truth and assertibility (in the specified sense) equivalent notions?

They are not. Assertibility is always relative to an epistemic context of evaluation. In most cases truth is not. Some philosophers have discussed the application of a version of truth-relativism to special cases, like sentences expressing judgments of taste (see Carpintero and Kölbel 2008). But for most sentences, for example for token-sentences such as “my father is Polybus” or “the equation has no solutions, if n is greater than 2”, truth (or falsity) is not relative to a context of evaluation.

Assertibility is relative to a context of evaluation because justifiedness is relative to such a context. When the boy Oedipus in Corinth, raised as the son of Polybus, says “My father is Polybus king of Corinth”, in that context of evaluation the assertion is justified: it meets all objections in the context. But Oedipus undertakes the commitment that the asserted sentence be *true*, not that it be simply assertible by virtue of a justification that is provisionally accepted in that context of evaluation. If the token sentence is true, it will remain true. If it is only assertible in a context of evaluation, it can cease to be assertible. This is precisely what happens in the story of Oedipus. Towards the end of the story, in the new context of evaluation in Thebes, Oedipus has to admit that his past assertion made in Corinth is not really justified. What in Corinth was a justification turns out to be misleading in light of the new information acquired in Thebes. A concatenation of cogent inferences compels Oedipus to accept the unpleasant truth about his origins: he is the son of Laius. Oedipus is a truth-seeker. That is why he cannot rest content with the old misleading evidence, which he once believed to amount to a justification. If he were not an intellectually virtuous investigator, he would be free to ignore the decisive objections that, in the new context of evaluation in Thebes, cannot be answered in a satisfactory way. But he is intellectually virtuous and, though it is painful to him, he is epistemically compelled to retract the unjustified assertion about Polybus’ paternity and to accept other assertions that he previously rejected. The story of Oedipus illustrates how justifiedness and assertibility are relative to the context of evaluation (for a similar view on justified belief, see Annis 1978).

But if we agree that justification is relative, we do not have to conclude that truth is relative. On the contrary, it is our commitment to a notion of non-relative truth which forces us to reject a misleading presumed justification and drives us beyond the limits of a given epistemic context. In general, even if an investigator finds an argument to support an assertion, so that the argument stands up to all the objections raised in the given context of evaluation and the uttered sentence is assertible relative to that context, the virtuous investigators will not necessarily stop investigating about the issue related to that assertion. Carefulness, curiosity, critical attitude or other intellectual virtues often urge them to investigate further. The investigators are guided by

the idea that though the assertion is justified with respect to a context-relative notion of justifiedness, it might still be untrue with respect to a non-relative notion of truth. The latter notion is more important, because in making the assertion the proponent has undertaken the commitment that the sentence is true. In a new epistemic context a mistake can be detected; the old argument can thus be rejected, and the assertion can be retracted. Thus in the new context the asserted sentence can cease to be assertible. By contrast, if it is true, starting from a certain stage of the investigation the sentence will remain assertible. Intellectually virtuous investigators seek truth, and not simply assertibility in a context. They do not rest content with a given context of evaluation. They articulate and enrich the contexts of evaluation, because they are driven by their commitment to truth. The truth-commitment is not only the commitment that the asserted sentence be assertible in a context of evaluation C , but that it remain assertible in other more articulated contexts generated by further investigation. The truth-commitment goes beyond the given epistemic context. So the act of assertion involves two normative notions: justifiedness (or assertibility), which is relative to a context of evaluation, and truth, which is transcontextual.

We can answer the question concerning the act of inference in a similar way: is cogency relative to a context of evaluation the only norm that distinguishes good inferences from bad ones? Justifications are mostly arguments composed of cogent inferences. The relativity of justification and the relativity of cogency are obviously connected. But if assertions are subject to two normative notions, relative justifiedness and trans-contextual (cross-contextual) truth, we may surmise that inference too is subject not only to a norm of relative cogency, but also to a trans-contextual notion of validity. According to our analysis of inference, a proponent who makes an inference implicitly undertakes the commitment that the premises rationally support the conclusion. What kind of commitment is this? Is it the commitment that the inference is cogent in a given epistemic context, perhaps in the context of utterance? If we now think of deductive inferences, it seems clear that the answer is in the negative. In deductive inferences the commitment is that the premises support the conclusion *conclusively*. We have seen that the truth-commitment of an assertion goes beyond the present epistemic context. In a similar way the conclusive support-commitment goes beyond the present epistemic context. If participants in an epistemic context C_i make a conclusive inference, they aim beyond cogency in C_i . The investigators do not simply aim at inferences that hold good in C_i ; they aim at inferences that stand up to criticisms and continue to be cogent even in the new contexts generated by C_i ; they aim at *stable cogency*. Leibniz himself would now say that his aforementioned inference in *Nouveaux essais* was wrong, because

that inferential act involved a support-commitment that was not merely a commitment to cogency in the context of utterance, but a commitment to stable cogency. This suggests the following definition of *validity* of an inference. A *valid inference* in an epistemic context C_i is an inference that is cogent in C_i and continues to be cogent in all new epistemic contexts C_{i+1}, \dots, C_n that are generated by C_i .

Now we can go back to the question: is the only available normative notion that distinguishes good and bad inferences the notion of cogency relative to a context of evaluation? The answer is: no, cogency relative to a context is not the only relevant normative notion. We have another trans-contextual notion, the notion of validity, which is stronger and more objective than relative cogency. But our knowledge of validity is fallible. We can never be certain that an inference is valid. In some future epistemic context someone might advance a successful objection.

Acknowledgements I am grateful to Bernhard Weiss and to the participants in the workshop “Inferences and Proofs” in Marseille for helpful comments on earlier versions of this paper.

References

- Annis DB (1978) A contextualist theory of epistemic justification. *Am Philos Q* 15(3):213–219
- Aristotle (2002) *Nicomachean ethics*, translated by C. Rowe, introduction and commentary by S. Broadie. Oxford University Press, Oxford
- Axtell G (ed) (2000) *Knowledge, belief and character*. Rowman & Littlefield, Lanham
- Baehr J (2011) *The inquiring mind*. Oxford University Press, Oxford
- Boghossian P (1997) Analyticity. In: Hale B, Wright C. (eds) *A companion to the philosophy of language*. Blackwell, Oxford
- Carpintero MG, Kölbel M (eds) (2008) *Relative truth*. Oxford University Press, Oxford
- Cellucci C (2009) The universal generalization problem. *Logique et Anal* 52(205):3–20
- Code L (1984) Toward a ‘Responsibilist’ epistemology. *Philos Phenomenol Res* 45(1):29–50
- Cozzo C (1994) *Meaning and argument*. Almqvist & Wiksell, Stockholm
- Cozzo C (1998) Epistemic Truth and Excluded Middle. *Theoria* LXIV(2–3):243–282
- Descartes R (1985) Rules for the direction of the mind. In: Cottingham J, Stoothoff R, Murdoch D. (eds) *The philosophical writings of descartes*, vol I. Cambridge University Press, Cambridge
- Dummett M (1978) *Truth and other enigmas*. Duckworth, London
- Dummett M (1981) *Frege. Philosophy of language*, (2nd edn). Duckworth, London
- Fairweather A, Zagzebski L (eds) (2001) *Virtue epistemology, essays in epistemic virtue and responsibility*. Oxford University Press, Oxford
- Frederick S (2005) Cognitive reflection and decision making. *J Econ Perspect* 19(4):25–42
- Frege G (1884) *Die Grundlagen der Arithmetik*. Verlag von Wilhelm Koenner, Breslau. English translation by J. L. Austin (1959) *The foundations of arithmetic*. Blackwell, Oxford

- Gorgias (2005) *Encomium of Helen* (ed. by D. M. MacDowell). Bristol Classical-Duckworth, London
- Kahneman D (2012) *Thinking, fast and slow*. Penguin, London
- Leibniz G (1990) *Nouveaux essais sur l'entendement humain*. Flammarion, Paris
- Mares ED (2004) *Relevant logic. A philosophical interpretation*. Cambridge University Press, Cambridge
- Pears D (2008) *Paradox and platitude in Wittgenstein's philosophy*. Clarendon Press, Oxford
- Plato (1997) *Complete works*, ed. by J. M. Cooper. Hackett, Indianapolis-Cambridge
- Prawitz D (1973) *Towards a foundation of a general proof theory*. In: Suppes P et al (eds) *Logic methodology and philosophy of science*, vol IV. North Holland, Amsterdam
- Prawitz D (2005) *Logical consequence from a constructivist point of view*. In: Shapiro S (ed) *The oxford handbook of philosophy of mathematics and logic*. Oxford University Press, Oxford
- Read S (1988) *Relevant logic*. Blackwell, Oxford
- Roberts CR, Wood WJ (2007) *Intellectual virtues. An essay in regulative epistemology*. Clarendon Press, Oxford
- Thurston W (1998) *On proofs and progress in mathematics*. In: Tymoczko T (ed) *New directions in the philosophy of mathematics. Rev. and Exp. Ed.* Princeton University Press, Princeton
- Williamson T (2017) *Semantic paradoxes and abductive methodology*. In: Armour-Garb B (ed) *Reflections on the liar*, Oxford University Press, Oxford
- Wittgenstein L (1956) *Bemerkungen über die Grundlagen der Mathematik (Remarks on the foundations of mathematics)*, ed. by G.H. von Wright, R. Rhees, G.E.M. Anscombe. Blackwell, Oxford