

## THE STRENGTH OF ARGUMENTS

Eivind Kolflaath

Department of Philosophy and Faculty of Law, University of Bergen

The Toulmin scheme is a tool for analyzing the *structure* of arguments. In this talk I will discuss whether the Toulmin scheme includes all the elements that are necessary for assessing the *strength* of an argument.

First, I want to point out that the notion of strength should not be defined in terms of validity. By “valid argument” I mean a set of premises related to a conclusion in such a way that it would be irrational to doubt the conclusion if we accept all the premises. It is tempting, perhaps, to think of valid arguments as maximally strong arguments, and of strong inductive inferences as arguments that are “almost” valid and thereby *quite* strong, though not as strong as a valid argument. In that case, we are taking the notion of validity as the ideal for all kinds of arguments. The closer the argument comes to this ideal, the stronger it is.

Valid arguments, however, are first and foremost *safe* arguments. Some are quite trivial and do not deserve our attention. For example, it is completely safe to infer a disjunction from one of its members. In general, validity is just a matter of combining sufficiently strong premises with a sufficiently weak conclusion. In a valid argument, the conclusion is not stronger than the premises taken together. Thus, the stronger premises, the more likely is the argument to be valid. Certainly, in real life we require that all the premises of an argument are true, or at least appear to be true. In other words, the premises cannot be *too* strong. Yet, the fact remains that validity is partly a matter of having a sufficiently weak conclusion. (What would be a sufficiently weak conclusion depends, of course, on the set of premises.) One way of weakening the conclusion is to add some kind of reservation to the claim.

Perhaps some examples would be appropriate. The following pair of arguments, or at least something similar to it, has been used quite often to help students understand the difference between a valid and an invalid argument:

- (a)     (1) If it has been raining recently, then the streets are wet  
          (2) It has been raining recently  
          (3) The streets are wet

- (b) (1) If it has been raining recently, then the streets are wet  
(2) The streets are wet  
(3) It has been raining recently

What we tell our students is that while (a) is a valid argument, the conclusion in (b) could be false even if both premises were true. Thus, the reasoning in (b) is unjustified and simply irrational – at least that is what some of us tell our students.

Maybe (a) and (b) together is an instructive example of the difference between valid and invalid arguments. But there is also something quite misleading about it. What we usually don't tell our students is that (b) is *not* the way we actually reason when we look out the window and notice that the streets are wet. Rather, our way of reasoning would be more like the following:

- (c) (1) The streets are wet  
(2) Almost always, if the streets are wet, it has been raining recently  
(3) There are no indications that it hasn't been raining recently  
(4) Presumably, it has been raining recently

This is a perfectly rational argument, and it is definitely valid in the sense mentioned above. It is valid partly because the order of cause and effect has been reversed in the general premise, but also because the conclusion has been weakened by adding a reservation so that the conclusion matches the strength of the premises.

What should also be noticed here is that (c), unlike (a) and (b), refers to some *epistemic situation* – that is, to the information available to the reasoning agent at the time the reasoning took place. Of course, the epistemic situation could change – maybe at some later stage the reasoning agent would find indications that it hadn't been raining after all. In that case, line (3) is not true and the conclusion is no longer justified.

The possibility that new information could alter the conclusion does not imply that (c) is an irrational argument. We make presumptive inferences all the time – in everyday life, within the sciences and in the courtroom – and it would be highly irrational *not* to do so. In fact, we wouldn't survive long if we didn't base our decisions on such inferences.

Now, the following argument is a representation of a famous example given by Toulmin in *The Uses of Argument* (p. 102). I know that there are people in the audience who disapprove that warrants are referred to as “premises”, but I do this only for the sake of conformity in my discussion.

- (d) (1) Harry was born in Bermuda
- (2) A man born in Bermuda will generally be a British subject
- (3) There is no information indicating that both his parents were aliens
- (4) There is no information indicating that he has changed his nationality<sup>1</sup>
- (5) There are no possible exceptions to (2) other than those mentioned in (3) and (4)
- (6) Presumably, Harry is a British subject

While (1) is the *data* (or datum) in Toulmin’s example, (2) is the *warrant*. In (3) and (4), we find two *exceptions* to the warrant. In (6), “presumably” is the (modal) *qualifier* while the rest of the sentence expresses the *claim* of the argument. I take it that (5) is part of Hitchcock’s version but not of the original version.

In the Toulmin scheme, exceptions are thought of as exceptions to the applicability of the warrant. In my abstract to this talk I wrote that exceptions are “ways in which the claim would not be true”. This is not correct, however, since in some cases the exception is consistent with the claim. In the above example, Harry could be a British subject even if both his parents were in fact aliens. The second exception, however, is not consistent with the claim – at least if we make the presumption that dual citizenship is not permitted. In that case, having changed his nationality is *an instance of the negation* of being a British subject.

In the textbook *An Introduction to Reasoning*, Toulmin et al. tell us that an inference “is warranted – [...] the claim is directly supported by the grounds – only *in the absence of some particular exceptional condition*, which would undercut (i.e., withdraw the authority of the warrant for) the inference” (p. 96). The strength of an argument depends in part, I think, on how many possible exceptions there are, to what extent they are specified and how easy it would be to decide whether or not an exception applies in this particular case.

---

<sup>1</sup> This is taken from Hitchcock’s version of the argument, cf. David Hitchcock: “Toulmin’s Warrants”, p. 1). Line (5) would not be justified in the original version of the example.

In the first example following the words just quoted, we have an argument with the datum that Hannah Smith is a local taxpayer, the warrant that all local taxpayers are normally entitled to vote at town meetings, and the claim that Hannah Smith is entitled to vote at town meetings (p. 96-97). The exceptions to the warrant include the possibilities that Hannah Smith is a noncitizen, a minor, a lunatic, or other disqualified person.

The authors comment that the warrant in this example “may hold good as a *general rule*, but it does so only subject to certain *specific exceptions and exclusions*”.<sup>2</sup> The last exception, however, is not very specific. In fact, with the addition of “or other disqualified person” the whole package of exceptions is nothing less than the entire negation of the claim, since “not entitled”, in this context, would simply mean the same thing as “disqualified”. Again, the argument is valid, but its validity has much in common with the truth of a statement like

“It was the defendant who committed the crime, unless he didn’t do it”

– that is, a statement of the form P or not-P. The argument is *safe* rather than strong. Certainly, the argument is better off than the statement I just mentioned since after all, even if we do not know what all the possible exceptions are, we do know that the warrant applies in most cases, whereas the statement does not tell us anything about which of P and not-P is the most probable. Still, an open-ended package of exceptions creates a kind of opacity. If we don’t know what all the possible exceptions are, we do not know what data we should look for in order to figure out whether one of the possible exceptions actually applies in this particular case. Thus, the argument we have just considered is less transparent than the argument in (d).

In many contexts, however, there is a more problematic kind of opacity – one which is not reflected by the Toulmin scheme. In real life the epistemic situation is quite often far more complex than in the examples above. If we want to apply the Toulmin scheme in, say, the O. J. Simpson case, in order to figure out what reasons there are for believing that Simpson really killed his ex-wife and her friend, we have to deal with an enormous amount of data, which would, I guess, require not just one, but several warrants that could link the various parts of

---

<sup>2</sup> Toulmin et al., p. 98.

the data to the claim that O. J. Simpson did commit these murders. In any case, what we know is that there are some data that increase the probability of the claim. But although we know *in what direction* these data change the probability of the claim, we do not know *to what extent* they do so. In such cases, it is often hard to tell whether the claim *is* probable at all – that is, more probable than its negation.

Toulmin points out in *The Uses of Argument* that “the adverb “probably” serves us as a means of qualifying conclusions and assertions, so as to indicate that the statement is made something less than positively, and must not be taken as committing the speaker to more than a certain extent” (p. 89-90). Of course, I agree that by saying that something is “probable”, one makes a weaker commitment than by saying that this something is certain. But on the other hand, by saying that something is “probable”, one makes a stronger commitment than by just saying that it is possible. By saying that something is “probable”, we state that this something is more probable than its negation.

Some would maintain that in cases like the O. J. Simpson example we should use words like “presumably” or “plausibly” rather than “probably”. In the textbook by Toulmin et al., the authors tell us that “there are two distinct reasons why claims often have to be presented as somewhat less than definite or certain” (p. 95). *Either*, it may be that the premises (which is not the term used by the authors) “lend their support to C only partly or weakly” (p. 96). *Or*, it may be that the premises “support the claim only in certain conditions” (p. 96). In the first case, the claim is *probably* true, while in the second case the claim is *presumably* true.

Maybe such a distinction could be maintained in concise examples such as those presented in this textbook and in *The Uses of Argument*, but in the more complex cases *both* phenomena are typically at play. In such a case, should we say “probably” or “presumably”? In Norwegian we tend to use the word “sannsynligvis” in both cases, although sometimes we say that something is “plausibelt” (i.e., plausible) or we state what is “tilsynelatende” (i.e., apparently) the case.

Now, even if we reject the idea that the notion of strength of arguments should be related to the notion of validity – or to the notion of *soundness*, which is validity together with true premises – it is still tempting to think that the strength of an argument depends on the degree of probability or plausibility of the claim given the data, warrant etc. supporting it. But this is

not an attractive notion of strength, since the strength of an argument would then be partly a matter of leaving out data that would decrease the probability of the claim. At least some of the texts analyzed by students at the Academic writing course here at the University of Bergen are heavily biased in this way – these texts present one-sided arguments rather than balanced discussions. I prefer a notion of strength that doesn't favour such arguments.

I think that the notion of strength should reflect not only the probability of the claim, but also its *robustness*. The more data the reasoning agent text takes into consideration – *both* data supporting the claim and data supporting its negation – the less vulnerable is the conclusion towards new information. This is the dimension that Keynes referred to as the “weight” of an argument,<sup>3</sup> and it is not the same thing as the probability dimension. A high probability could have low robustness, and vice versa.

Certainly, it's a rather trivial point that the conclusion we make is more robust the more relevant information we have taken into consideration. It is also trivial that the relevant information includes data that support the negation of the claim. But as reasoning agents we seem to forget this all the time. It's so tempting to accept a claim on the basis only of the reasons supporting it – we are so convinced that we think, quite often mistakenly, that we don't need to hear the reasons supporting its negation. The Toulmin scheme, focusing at least primarily on the reasons supporting the claim, does not remind us of this danger.

Of course, when applied in the more complex cases the Toulmin scheme could include data that do not support the claim, and even data that support its negation. Even so, the fact that the claim is the only hypothesis in the scheme makes the scheme biased towards that hypothesis.

Besides, in many contexts it would not be sufficient to compare that which supports the claim with that which supports its unspecified negation. In such cases, we need to consider instead various *instances* of this negation. For example, the unspecified negation of the claim that Napoleon was murdered by poisoning is that Napoleon was *not* murdered by poisoning (which includes, by the way, the possibility that Napoleon is still alive). An example of a particular instance of this negation, on the other hand, is that Napoleon died of cancer. We know that Napoleon is dead, and we know that he died of *something*. If I want to reach a

---

<sup>3</sup> Keynes: A Treatise of Probability, p. 71-78.

qualified opinion as to whether Napoleon was murdered by poisoning, I should look for data supporting this hypothesis and then do the same for each of the other hypotheses about Napoleon's death that seems to have at least *something* to it. If it turns out that the hypothesis that Napoleon was murdered by poisoning appears to be better off than each of the other hypotheses, then I conclude that Napoleon was probably murdered by poisoning. This is a *relative* or *comparative* notion of strength, but I think that in a context like this, that is all we are left with. The Toulmin scheme, on the other hand, does not reflect the relative notion but rather seems to presuppose some absolute notion of strength.

The need for comparing the claim with particular instances of its negation arises not only from the fact that a hypothesis tells us what data we should to look for, but also from the fact that quite often the data we have before us have to be interpreted. The interpretation of the data is typically governed by a hypothesis. Thus, we need competing hypotheses in order to shed *different* lights on the same data. In many contexts, the unspecified negation of the claim is too broad and opaque to serve as a hypothesis.

What I suggest, then, is that in the more complex situations we need an argument that includes a comparison of several hypotheses. To be more specific, what we need in such contexts is the kind of argument commonly referred to as "abduction" – that is, an *inference to the best explanation*. The broad structure of an inference to the best explanation can be outlined as follows:

*F* is a finding or given set of facts

*E* is a satisfactory explanation of *F*

No alternative explanation *E'* given so far is as satisfactory as *E*

Therefore, *E* is plausible, as a hypothesis

(Walton: *Legal Argumentation and Evidence*, p. 44)

*D* is a collection of data (facts, observations, givens)

*H* explains *D* (would, if true, explain *D*)

No other hypothesis explains *D* as well as *H* does

Therefore, *H* is probably correct

(Josephson: “Smart Inductive Generalizations are Abductions”, p. 31)

There are some minor differences between the two schemes, but these are not significant to my discussion. The point I want to emphasize, is that in both schemes the third premise takes care of the comparative dimension, and that this premise is thought of as necessary for the conclusion.

I am not saying that the Toulmin scheme is ill-suited as a tool in the analysis of the *structure* of the arguments presented in the texts that are analyzed as part of the Academic writing course. On the contrary, the absence of the comparative element fits the biased character of some of these texts. But the students should not be asked to judge the probability or plausibility of the claim on the basis of the text. That would be like asking a judge to make his or her verdict after hearing only one of the parties in the courtroom.