

RULES OF INFERENCE, OR SUPPRESSED PREMISES?

Originally published in *Mind* LXXIII (289), Pp. 84-96, 1964

[Aaron Sloman](#)

Oxford University

(Now at the [University of Birmingham School of Computer Science](#))

1. In ordinary discourse we often use or accept as valid, arguments of the form "P, so Q", or "P, therefore Q", or "Q, because P" where the validity of the inference from P to Q is not merely logical: the statement of the form "If P then Q" is not a *logical* truth, even if it is true. Inductive inferences and inferences made in the course of moral arguments provide illustrations of this. Philosophers, concerned about the justification for such reasoning, have recently debated whether the validity of these inferences depends on special rules of inference which are not merely logical rules, or on suppressed premisses which, when added to the explicit premisses, yield an argument in which the inference is logically, that is deductively, valid. In a recent contribution to MIND ("Rules of Inference in Moral Reasoning", July 1961), Nelson Pike describes such a debate concerning the nature of moral reasoning. Hare claims that certain moral arguments involve suppressed deductive premisses, whereas Toulmin analyses them in terms of special rules of inference, peculiar to the discourse of morality. Pike concludes that the main points so far made on either side of the dispute are "quite ineffective" (p. 391), and suggests that the problem itself is to blame, since the reasoning of the "ordinary moralist" is too rough and ready for fine logical distinctions to apply (pp. 398-399). In this paper an attempt will be made to take his discussion still further and explain in more detail *why* arguments in favour of either rules of inference or suppressed premisses must be ineffective. It appears that the root of the trouble has nothing to do with moral reasoning specifically, but arises out of a general temptation to apply to meaningful discourse a distinction which makes sense only in connection with purely formal calculi.

2. If A, B, C, . . . P, Q, . . . T are well-formed formulae in a formal system, and the sequence "A, B, C, . . . P, Q, . . . T" constitutes a "proof" in this system — a proof of T — then the question "Does the transition from P to Q depend on some other formula which is either an axiom or a predecessor of P in the sequence, or is it justified simply by a rule of inference?" is a clear question which will usually have a single correct answer, provided that a definite formal system is specified. This is because the formal systems in question are defined by means of their axioms and rules of inference, which must be specified either by means of a list, or recursively, so that if we know which formal system is referred to we can effectively decide whether the transition from P to Q accords with one of the rules of inference of that system or not. But different systems, containing similar symbols and well-formed formulae, while differing in their axioms and rules of inference, may nevertheless yield the same class of valid derivations and theorems. Even within one system it is possible for a step from one formula to another to be validated both by one of the rules of inference and also by another rule of inference together with a "suppressed premiss" which is a theorem or axiom. These facts alone should make us suspicious

of any attempt to make an absolute distinction between inferences which depend solely on rules of inference and those which require an additional premiss as well. The grip of this dichotomy on our imagination might be further loosened if we thought of formal systems as defined not in terms of their axioms and rules of inference, but in terms of the class of their well-formed formulae which are theorems, or in terms of the set of possible interpretations of the system. (The latter would be preferable, as it would focus our attention on questions of semantics which cannot be answered by the methods of symbolic logic, but which nevertheless deserve the attention of logicians.) I shall presently offer further reasons for not thinking of the distinction as important and absolute outside of the science of formal systems.

3. It will help to consider some examples of non-formal arguments, of which the first is Pike's example (p. 392) of an "ordinary moralist" called Seeforth who appeals directly to the fact that

- (1) stealing usually results in more suffering than not stealing, in order to justify his general moral principle that
- (2) men ought never to steal.

Two analyses of this argument of Seeforth's, neither of which questions the validity of the reasoning, are offered in order to explain the support which (1) provides for (2). According to the first analysis (Toulmin's), the inference is not deductive in nature, but depends on a special rule of inference, which may be described as "the rule of least suffering" and formulated thus

- (3) From a proposition having the form 'actions of type X usually result in more suffering than actions of type not-X' one may conclude a proposition of the form 'one ought never to perform actions of type X'. (op. cit. p. 393.)

This is supposed to be a special type of inference peculiar to the discourse of morality. (3) is not a *premiss* from which, together with (1), (2) may be deduced. It is supposed to be a rule which directly justifies the inference from (1) to (2).

According to the second analysis (following Hare), Seeforth has made a purely deductive inference using a suppressed premiss, e.g.

- (4) If actions of type X usually result in more suffering than actions of type not-X, then one ought never to perform actions of type X. (op. cit. p. 393.)

This suppressed premiss, if taken together with (1) would *logically* entail Seeforth's conclusion without the help of any special rule of inference peculiar to morals. This analysis is based on the view that there are no such rules of inference as (3) which validate arguments from (true) factual premisses to moral principles. All arguments must be based on premisses, moral and non-moral, which logically (e.g. syllogistically) support their conclusions.

4. The debate on the relative merits of these two types of analysis is described in Pike's article, and the arguments found to be inconclusive. But how are we to understand the question : Which analysis is correct? Various interpretations are possible, but before we turn to them, we must take note of the fact that if there is a problem it is not peculiar to moral reasoning, since wherever an argument may be justified by either an appeal to a rule of inference or an appeal to a suppressed premiss, it may also be justified by the other type of appeal. More accurately, if there is a rule of inference which permits Q to be inferred from P then there is a suppressed premiss S such that S and Q together logically entail P, and *vice versa*, and in either case all those inferences which are validated either by the rule or by the addition of the premiss S are also validated by both. This may be illustrated by the following examples. In each case both a rule of inference and a suppressed premiss are given.

(1a) Not all cupboards are bare, therefore (Premiss)

(2a) Some cupboards are not bare. (Conclusion)

This inference is validated by either of the following:

(3a) From a proposition having the form 'Not all cupboards are of type X' one may infer a proposition of the form 'Some cupboards are not of type X'. (The rule of inference.)

(4a) If not all cupboards are of type X then some cupboards are not of type X. (The suppressed premiss.)

Another example is the inference from

(1b) John is in the house,

to

(2b) John is upstairs or downstairs.

And this may be validated by appeal to either of the following:

(3b) From a proposition having the form 'X is in the house' one may infer a proposition of the form 'X is upstairs or downstairs'.

(4b) If a person is in the house, then that person is upstairs or downstairs.

Note that, as before, the rules of inference (3a) and (3b) are not supposed to be suppressed premisses which, when taken with (1a) and (1b) respectively, render the arguments deductively valid. They are supposed to be rules of inference which directly justify their respective arguments, given the truth of the premisses. Each is peculiar to a special type of discourse, though in the first case we might have offered instead of (3a) and (4a) the following topic-neutral rule and premiss: "From a proposition having the form 'Not all A's are of type X' one may infer a proposition of the form 'Some A's are not of type X'", and "If not all the things of a certain kind are of a certain type, then some things of that kind are not of that type". In either

case, the suppressed premiss together with the explicit premiss logically entails the conclusion, so that according to the analysis in terms of suppressed premisses there is no need for special nonlogical rules of inference. It should be clear that in connection with almost any kind of argument, whether it is concerned with morals or not, both types of analyses are possible, and if the question "Which analysis is correct?" arises in one case it arises in all cases.

5. We can see in general how to get from one kind of analysis to the other first of all by noticing that if the inference from a proposition P to another Q is in accord with a rule R, then all inferences from a single premiss to a conclusion which accord with R must do so in virtue of some common structural relationship between the premiss and the conclusion, in virtue of which they are instances of the application of the rule. Once we have discovered what this structural relationship is, we can, by translating from the formal to the material mode, formulate a suppressed premiss S, "If such and such a type of thing is the case, then so-and-so is the case", such that S and P together logically entail Q. This is illustrated by the examples already given. Secondly, and conversely, if we are told that the inference from P to Q is based on a suppressed premiss S, so that S and P together logically entail Q, then this entailment must hold in virtue of some formal relationship between S, P and Q. Now, if we consider all pairs of propositions F and G which stand in this relation to S, so that S and F together logically entail G, then we must find that there is some relation, not necessarily formal (see the last illustration, for example) between F and G common to all such pairs, in virtue of which the entailment holds. If we formulate a rule of inference which allows us to proceed from any proposition F to another G, if there is one, to which it stands in this last relation, then this rule will validate all those inferences which are logically validated by the addition of the premiss S, including our original inference from P to Q.

It is clear that this method of deriving a rule of inference from a suppressed premiss, or *vice versa*, so that any argument is validated by the rule if and only if it is logically validated by the addition of the suppressed premiss, is of quite general application. So any analysis of a particular argument in terms of a suppressed premiss may be replaced by an equally general analysis in terms of a rule of inference and *vice versa*. This is not a peculiarity of moral arguments, nor of "rough and ready arguments", as suggested by Pike on page 398.

6. We must now return to the question: Which analysis is correct? It is not at all clear what is meant by this question. How are we supposed to decide which kind of analysis is correct? What factors are relevant? Perhaps the arguments on both sides were found by Pike to be ineffective on account of the indeterminacy of the *question*. I shall offer several interpretations of the dispute and try to show that on each interpretation either both sides are correct, or neither, except in so far as the dispute is a terminological one. Suppose we return to Seeforth's inference from (1) Stealing usually results in more suffering than not stealing, to (2) Men ought never to steal, and consider the two analyses set out in §3, above, in terms of a special rule of inference, "the rule of least suffering", and in terms of a suppressed moral premiss. Then to say that one of the analyses is correct might mean (a) that it describes what *Seeforth* really meant when he offered (1) as a justification for his assertion of (2), or what is *usually* meant by anyone who uses such an argument; or (b) that the analysis shows what Seeforth *ought* to have meant, that is what he

would have meant had he been free from confusions and clear about what he was doing; or (c) that the analysis justifies his argument in a way in which the other analysis does not; or, finally, the assertion that one analysis is correct and the other incorrect might mean (d) that one of them describes the inference accurately whereas the other involves a misuse of logical terms. I shall discuss these interpretations separately.

(a) It is clear that it is not for philosophers to decide which analysis is correct, if either is, in the first sense, for it is surely an empirical question whether Seeforth meant his argument to be based on a suppressed premiss or a special rule of inference. We may be able to find out by asking him. On the other hand, we may be unable to find any evidence that he means one thing or another, since people are able to argue, and to distinguish good from bad arguments without even having the concepts "suppressed premiss" and "rule of inference", in which case they can hardly *really* intend their arguments to be validated by one rather than another. We might try to discover his intentions by asking him the two questions "Did you offer this as a reason because you believed that so-and-so is the case?" and "Did you offer this as a reason because one can always go from a statement like this to a statement like that?" But then we must not be surprised to get positive answers to *both* questions, or to *neither*.

I conclude that in general neither type of analysis would serve any better than the other to explain what was really meant by someone who offered such an argument, and if either did turn out to be a correct analysis in this sense on some occasion, this would be merely a contingent fact to be discovered empirically and not by philosophical argument. Pike does not seem to interpret the question in this first sense ("It is not . . . a dispute about the explicit content of Seeforth's reflections" -- p. 393), but only if it *were* a dispute of this sort could the ineffectiveness of the arguments be explained in terms of the "rough and ready" nature or "shaggy and wrinkled hide" of actual moral reasoning (pp. 398-399). This should become clear in the discussion of other interpretations of the dispute.

(b) Is it possible for only one of the analyses to be correct in the second sense, namely by virtue of stating what Seeforth *ought* to have meant? But what mistake could be made by one who intended his argument to be based on the rule of inference, or on the suppressed premiss? If it is true that his conclusion may be inferred from his premiss in accordance with his rule, then it is also true that adding the suppressed premiss would make his argument deductively valid, and conversely. Neither explanation of his argument, if offered by Seeforth himself, would show that he was more confused or less intelligent than if he had offered the other. There is no more error or confusion in one expanded version of his argument than in the other. Of course, there is always the further question: How does the explanation help to establish Seeforth's right to assert the conclusion? But this question arises for both analyses, in the form "Why are inferences made according to this rule valid?" or "Why is this premiss true?" If an answer can be given to either of these questions it may be transformed into an answer to the other, as will be shown presently. I agree, therefore, with the conclusion of Pike's discussion on pages 394-398, namely that from this point of view there is nothing to choose between the two types of analysis.

7. (c) The third interpretation of the question "Which type of analysis is correct?" is a little more interesting. On this interpretation the question is whether there is some sense in which either appeal to a rule of inference or appeal to a suppressed premiss provides a better *justification* for Seeforth's argument, or for any of the other arguments which allow both types of appeal. A justification here presumably means an answer to a question like "Why does this follow from that?" or "Why, if that is true, must this be true?" Answers to this sort of question seem to fall into two main classes, depending on their function. The first kind of justification states that the inference is of a certain type, or exemplifies a certain pattern, and implies that no inferences of that type, with that pattern, lead from true premisses to false conclusions, or at least that it is reasonable to assert the conclusion of such an inference if one knows that the premiss is true. The second sort of justification points to something which guarantees or proves or at least makes it likely that no inferences of the type in question yield false or unreasonable conclusions from true premisses.

The first sort of justification may indicate a pattern in the inference by specifying some kind of relationship which holds between premiss and conclusion, whether syntactical, logical, or semantical. We are concerned with two rival methods of picking out a pattern of which the first attempts to justify the inference by pointing to a rule of which it is an application, implying that no inferences which accord with that rule, i.e. no inferences with the same pattern, have true premisses and false conclusions, or lead from premisses known to be true to conclusions which it is not reasonable to believe. The second justification points to a suppressed premiss which turns the inference into a logically valid one, implying that no inference which becomes logically valid on the addition of this premiss leads from true premisses to false or unreasonable conclusions. But if this sort of justification is required, then there is obviously nothing to choose between the rule and the premiss provided that they are selected by the method explained in §5, for in that case every inference which is justified by one of them is justified by the other. (See also the beginning of §4.)

In any case, no justification of this first sort would normally be considered complete unless it was accompanied by, or itself intended to be, a justification of the second sort, which not only specifies a pattern or type of acceptable inference, but also indicates something which *guarantees* that inferences of this type will not or cannot lead from a true premiss to a false conclusion, or at least guarantees that it is reasonable to assert the conclusion if one knows that the premiss is correct. But it seems that neither appeal to a suppressed premiss nor specification of a rule of inference can provide a complete answer of this sort to the question "Why does this follow from that?", or, at any rate, neither gives a better answer to the question. For simply to point to a premiss which one has taken for granted is not to show why that premiss is or must be true, or why it is reasonable to accept it as true for purposes of logical inference ; and to point to a rule of inference which one has followed is not to show why that rule will not, or cannot, lead one astray. If one plays games with symbols, then one may invent one's own rules or initial configurations, but if one intends to say something with a meaning then a further justification may be required. (See, in this connection, "The Runabout Inference-Ticket", by A. N. Prior in *Analysis*, December 1960.)

In either case, the justification, if there is one, must lie outside the mere pattern to be found in the inference. It may lie in some empirical fact, for example the empirical fact which justifies the inference from "This is a frog" to "This was once a creature without legs". Or it may lie in a linguistic convention, such as the one in virtue of which "Smith is a bachelor" says the same thing as "Smith is an unmarried man" and may therefore be inferred from it. Or it may lie in some feature of the semantic and syntactic rules according to which we use our logical words, such as the one which justifies the inference from "Not all cupboards are bare" to "If there are any cupboards then some are not bare". Or there may be a pragmatic justification for fixing the meaning of the word "true" in moral contexts in such a way that arguments with the same pattern as Seeforth's yield only "true" conclusions from true premisses.

So the complete justification for an argument cannot lie simply in the fact that it presupposes a certain premiss or in the fact that it proceeds according to some rule of inference. If there is something which justifies the use of a suppressed premiss, then we know (from §5) how to find a rule whose use is also justified, and which supports the same arguments. Similarly, if there is anything which justifies the use of a rule of inference, then we can find a suppressed premiss the use of which it would justify in the same contexts. Whatever it is that supplies the guarantee of the usefulness of the pattern of inference is no more closely connected with the suppressed premiss than with the rule. This is, of course, commonly denied, for example when philosophers argue that definitions can be used only to transform truths, not to found them. (See, for example, Quine's essay "Truth By Convention" in "Essays for A. N. Whitehead". Also Waismann's "Analytic-Synthetic" in *Analysis* (Dec. 1949).) But why should my decision to use the word "gleen" as a synonym for "glossy and green", i.e. my decision to use these expressions to refer to the same properties, be any less closely connected with the truth of "A thing is gleen if and only if it is green and glossy" than with the validity of the principle that substitution of "gleen" for "glossy and green", or *vice versa*, in certain contexts will not lead from true statements to false ones? Both are logical consequences of the logical relation holding between the meanings of the words "gleen", "glossy" and "green", and, of course, the fact that our logical constants have those functions which they do have. What could it possibly mean to say that either the premiss "A thing is gleen if and only if it is glossy and green" or the aforementioned principle provides a better or more fundamental justification for the inference from "This is gleen" to "This is glossy and green" than the other?

Of course, where the suppressed premiss is not an analytic truth, but some other sort, we are inclined to think of it as providing a better justification for an inference, since it states the *fact* which justifies inferences of that type. For example "Every frog was once a creature without legs" states the empirical fact which justifies the inference from "This is a frog" to "This was once a creature without legs". On the other hand, the principle of inference, such as "From a proposition of the form 'x is a frog' one may infer a proposition of the form 'x was once a creature without legs'", does not *state* that fact, and so it cannot be used to provide as good a justification. It might be argued against this, however, that the principle does state the same fact, but in another way. Or instead it might be argued that the suppressed premiss does not provide a better justification since it does not state that adding it to the explicit premiss makes the argument logically valid, whereas the rule of inference does say that the conclusion may be inferred. More simply, just as we

need to state a fact which justifies the employment of the rule of inference, BO we need to formulate a rule of inference which justifies the claim that the suppressed premiss supports the original argument. If a defender of suppressed premisses replies that he mentions such logical *principles* of inference in his analysis of the argument, his opponent may reply that he can mention the *facts* which provide a justification in his analysis. It should be clear that there is no real dispute here about the nature of the inference, only a dispute as to the best form of words with which to describe it. This is a dispute of the fourth kind, which must now be discussed.

8. (d) On the fourth and last interpretation, the question "Which type of analysis is better" is a terminological one, to be settled not by finding out what Seeforth, or anyone else, really means, nor by enquiring into the relative degrees of confusion and clarity in various processes of reasoning, nor by asking what it is that justifies the inference in question, but by considering whether one of the analyses fails to correctly describe what happens, as a result of misusing some logical term or other, such as "rule of inference" or "premiss". There is no standard terminology outside of the study of formal systems to which appeal can be made here, and there are different ways in, which the terminology which applies to formal systems may be extended so as to apply to real languages, in which the words have meanings, various factors favouring one or other way of extending the terminology. It is possible, therefore, that philosophers who debate the relative merits of the two types of analysis may be involved in a dispute as to the way in which philosophical terminology should be made precise and unambiguous, and fail to see that this is their only disagreement, as a result of their mistaken inclination to assimilate arguments used in real life to proofsequences in formal calculi. In this way they lose sight of non-terminological issues such as whether and how moral or inductive inferences may be said to be justifiable, while thinking that they are still debating these issues.

The following are examples of the sorts of considerations which are relevant to the terminological question. A defender of analyses like Hare's, in terms of suppressed premisses might argue that the use of the expression "rule of inference" should be restricted in two ways. First of all, it should not apply to anything which was not purely formal, i.e. topic-neutral, and secondly it should exclude any matters of substance. This would presumably leave only logical principles and those which depend on arbitrary definitions. The term "rule of inference" should apply only to topic neutral principles, first of all because this would help to exclude matters of substance, since they are usually concerned with some special kind of discourse, and secondly because in its original use in formal systems the term was used to correspond only to *patterns* of argument. That is to say, rules of inference were formulated only in terms of the logical form of premisses and conclusion, and to retain this convention would make theories of inference simpler and more elegant, at least in one respect, and more generally applicable. (This argument was suggested to me by R. H. Stoothoff.) Matters of substance should be excluded from the range of application of the expression "rule of inference" because this would force us to embody them in premisses, whether explicit or suppressed, and, since premisses more obviously invite the question "Why?" than rules of inference, this procedure would make it more difficult for philosophers, and others, to shirk questions of justification, or to give unsatisfactory answers by talking about "immediate connections" or "the rules of the language of morals". It is clear that the "rule of least suffering" (3) embodies a matter of substance, since one who

disagreed with it would not be expressing a mere terminological disagreement, but at least a practical one, and so these considerations would favour the adoption of Hare's terminology for describing Seeforth's argument.

Several arguments are available in support of a different terminology. First of all, it might be suggested that the word "premiss" should be applied only to explicit premisses, for it is not clear in what sense anything can be a premiss of an actual argument, when it is not only not made explicit in the argument, but not even thought of at the time by the person who uses the argument. After all, we can say which argument we are talking about only if we specify the premisses actually used. If we then add new premisses, we have a new argument. The suggestion is that the notion of a "suppressed premiss" is almost (but, of course, not quite) self-contradictory. A further difficulty is that it is not quite clear what we are supposed to allow to be a suppressed premiss in any particular argument. For if any one proposition is offered, then there will be many other propositions of different forms which are logically equivalent to the one suggested, and they too would turn the original argument into a logically valid one if they were added to it. Are we to say that they are all suppressed premisses of the argument? If not, what means are there for deciding which is the most suitable candidate? Or is the suppressed premiss some kind of "superpremiss" to the effect that all of these are true? But here we get the same trouble once again. Perhaps we are supposed to find out which is the right premiss by asking the person who offered the argument which of them he believes. But then whether a particular proposition is the suppressed premiss of an argument or not, will have to be settled by empirical investigation. What if he refuses to assent to any of them, owing to confusion or failure to perceive some logical connection, even though he stoutly maintains that his original inference is valid? Roughly the same sort of point could be made in a different way. We know which argument we are talking about because we know what the premisses are and what the conclusion states. But if an analysis is supposed to explain the reasoning involved in this argument, then it must show how the conclusion may be reached from the premisses, and anything which does this just is a rule of inference. This is what the expression "rule of inference" means, namely some procedure for deriving one statement from others. Hence even if we talk about suppressed premisses which turn the argument into a logically valid one, we are nevertheless going in a roundabout way to describe a rule of inference.

A further point might be made in connection with inferences whose justification is pragmatic, inferences which we describe as reasonable. In many such cases we should find it very unnatural to say that there was a suppressed premiss. For example, we regard it as reasonable for someone to argue: "All the cigarettes out of packets like this have contained tobacco in the past, so the cigarettes in this packet will contain tobacco." But we do not want to say that he believes that what has happened in the past will happen again, for we know quite well that he does not, and that he would not accept such a proposition as a deductive premiss. So the only alternative is to say that he adopts a rule of inference, and that this is pragmatically justified, even though it is capable of leading him astray, that is, leading him from true statements to false ones. Here it seems to be quite natural to speak of a matter of substance as embodied in a pragmatically justified rule of inference, which it is reasonable to follow, and only by some rather artificial convention could it be replaced by a premiss. These considerations would favour Toulmin's description of Seeforth's argument.

I do not think that there is any conclusive general argument in favour of one kind of terminology or the other. Whether either or both types of analyses should be accepted or not seems to be a problem to be settled by considering questions of convenience of description. Perhaps a whole new terminology should be devised. I cannot see that it is nearly as important a question as some of the other questions which I have mentioned, namely questions about the sense in which moral reasoning may be justified, if there is any. But even if there were a conclusive general argument in favour of one type of terminology, which therefore showed that one of the analyses involved a misuse of logical terminology, this would nevertheless not establish anything about what Seeforth's means, or ought to mean, nor would it show that a better justification for his argument is provided by a rule of inference than by a suppressed premiss or *vice versa*. All this should be clear from the discussion of the previous sections.

9. I conclude therefore, that in so far as the question "Which type of analysis is correct?" has an answer of interest to philosophers, it must be a question about the convenience of adopting a certain logical terminology, and not a question about the justification of moral reasoning or any other kind. It is because the question has not been given a clear meaning at all that the arguments so far brought forward on either side have been inconclusive. It is not, as Pike suggests, due to any peculiarity of moral reasoning, nor is it due to any "rough and ready" feature of informal reasoning. This conclusion would apply equally well to reasoning involving very precise concepts of Physics or Mathematics: to anything which is not just a proof-sequence of a formal calculus but an argument with meaningful premisses and conclusion. The generality of this conclusion seems to imply that certain philosophical debates about the role of scientific theories are in fact not debates about scientific reasoning at all, but about the relative merits of different systems of philosophical terminology.