

Are Some *Modus Ponens* Arguments Deductively Invalid?

DOUGLAS N. WALTON *University of Winnipeg*

Abstract: This article concerns the structure of defeasible arguments like: 'if Bob has red spots, Bob has the measles; Bob has red spots; therefore Bob has the measles.' The issue is whether such arguments have the form of *modus ponens* or not. Either way there is a problem. If they don't have the form of *modus ponens*, the common opinion to the contrary taught in leading logic textbooks is wrong. But if they do have the form of *modus ponens*, doubts are raised about the conventional dogma that all arguments having the form of *modus ponens* are deductively valid. By carefully examining arguments on both sides of the issue, reasonable doubts are raised about the view that all arguments having a *modus ponens* form are valid.

Resume: On concentre notre attention sur la structure des arguments tels que "Si Bob a des papules rouges, il a la rougeole; Bob a des papules rouges; donc il a la rougeole". Le point en litige est a savoir si de tels arguments ont la forme de *modus ponens*. S'ils n'ont pas cette forme, alors l'opinion courante exprimee dans les principaux manuels de logique est erronee. Mais s'ils ont la forme de *modus ponens*, on souleve des doutes au sujet du dogme conventionnel selon lequel tout argument ayant cette forme est valide. Un examen soigne des arguments opposes constituant cette controverse fait soulever des doutes raisonnables sur l'avis que tout argument ayant la forme de *modus ponens* est valide.

Keywords: argumentation schemes, defeasible inference, abductive inference, generalizations, artificial intelligence, conditionals, argument forms, diagnostic reasoning, legal argumentation, *modus ponens*.

1. Introduction

It is widely accepted in logic that arguments having the form *modus ponens* are deductively valid. You could almost say that this view is universally accepted. As shown below, it is taught in the leading logic textbooks. It is such a centrally presumed assumption in logic that any attempt to challenge its truth will almost certainly be rejected by any logician as highly implausible, or even misguided. What is argued below, however, is that there are many common arguments used daily in everyday reasoning that have the form *modus ponens* but are not deductively valid. Once this new view of the matter is developed, the currently prevailing view begins to seem more and more like a dogma. Which view is right? As shown below, neither view is trivial. Strong arguments and apparent proofs can be given for both views. The ultimate conclusion arrived at below is that the new view is

right. The arguments for the new view are mainly based on what are taken to be a number of counter-examples to the accepted view. These alleged counter-examples to the validity of *modus ponens* are not exotic or contrived. They are extremely common arguments of the most mundane kind. True, they have been sadly neglected by logic in the past, until the recent concerns about defeasible arguments have attained such a high visibility, especially in computer science (Reiter, 1985; Prakken, 1997). But they are plentiful, and they do, at least subject to the discussion below, arguably appear to have the form widely called *modus ponens*. To those who come to accept the new view, the prevailing view comes to be seen as a dogma of a past era.¹

What are the implications of accepting the new view and rejecting the currently accepted one? Does conceding that arguments of the *modus ponens* form can be deductively invalid in some cases undercut the whole basis of deductive logic? It is argued below that it does not. Instead, it has the effect of opening up new and important avenues of logical investigation that have been neglected, and even regarded as beyond serious investigation in the past. One outcome of exposing the dogma is to open up a new approach to the study of argumentation schemes.²

2. The Current View

An example of how current writing in philosophical logic adopted the view that arguments of the form of *modus ponens* are deductively valid is the following case from Sanford (1989, p. 16).

The Emil Argument

If Emil has been captured, then all is lost.
Emil has been captured.
Therefore, all is lost.

In discussing the Emil argument Sanford (p. 18) argues that it is a deductively valid argument. The reason he gives is that "it is impossible that the premises should all be true while the conclusion is false." (p. 18). But is it impossible that the premises should all be true and the conclusion false? I would say not. I would say that the premises being true gives a good reason for thinking that the conclusion is also true (other things being equal). But still, imagine that the premises are true but then unexpectedly, the enemy soldiers come down with a bad case of some disabling illness. Maybe, in such a case, all is not lost, despite Emil's capture. Such an event of mass illness among the enemy is not likely, but all it has to be is possible for the premises to come out true and the conclusion false. Since it is possible, I would say that the Emil argument above is not deductively valid. Of course you could take it to be deductively valid, as Sanford does. But what I am saying is that such an interpretation does not represent a realistic or accurate view of how the inference would be taken in everyday argumentation. Although such an inference would be taken as giving a reason to accept the conclusion as true, on the assump-

tion that the premises are true, it would not be taken as saying that if you accept the premises, you (necessarily) have to accept the conclusion.

But of course, what the committed formal logician (CFL) will do is appeal to something called the TF principle in section six below: for any conditional, if the antecedent is true and the consequent false, the whole conditional must be false. Using this principle, the CFL will argue that if the premise 'Emil is captured' is true, and the conclusion of the Emil argument turns out to be false, then the conditional "If Emil is captured, then all is lost" cannot really be true. Suppose the enemy had a mass illness, and the battle was won, even though Emil had been captured by the enemy before the battle. Then the CFL will say that the statement that if Emil is captured all is lost simply cannot be true. Therefore, the story goes, *modus ponens* is not invalid after all, because in the Emil case, the argument does not go from true premises to a false conclusion. This objection, which could be called the TF objection, takes us into a consideration of conditionals, and under what conditions they are defeated, in section six below. I will argue that the TF objection can be answered.

You don't have to look very far through current logic textbooks to find comparable examples of arguments like the Emil argument that are identified as having the *modus ponens* form. Consider the following example of an argument treated in (Layman, 1999, pp. 30-31).

If it is raining, then the ground is wet.
It is raining.
So, the ground is wet.

According to Layman (p. 31), this *modus ponens* argument is "obviously valid." And so it is, if you take the first premise to be false if the antecedent is true and the consequent false. But of course, it could be true that it is raining even though it is false that the ground is wet, in a given instance. For example, the ground could be covered. So questions can be raised whether the example really is a *modus ponens* type of argument, whether it really is a deductive argument, and whether it really is deductively valid.

Or consider the following example of an argument treated by (Hurley, 2000, p. 356).

If accidents have occurred in nuclear power, then we can expect accidents in nuclear weaponry.
Accidents have occurred in nuclear power.
Therefore, we can expect accidents in nuclear weaponry.

According to Hurley (p. 356), this argument has the *modus ponens* form and is deductively valid. But the same sort of questions can be raised about this example.

The general policy advocated by logic textbooks is to consider arguments like those above as *modus ponens* inferences that are deductively valid. Those of us who are accustomed to dealing with defeasible arguments as a distinctive type in their own right may find this practice questionable. But why have the logic textbooks persisted for so long in routinely classifying these defeasible arguments as

being straightforwardly deductively valid? The answer seems to reside in assumptions about how to translate conditionals from natural language argumentation into formal logic. These assumptions are indicated by the way the logic textbooks tell the student to translate conditionals by using the truth-functional hook (horseshoe). For example, Copi and Cohen (1998, p. 363) tell their readers that the statement, "If he has a good lawyer, then he will be acquitted." should be symbolized using the hook (the symbol for the material conditional) as the connective. By this symbolization, the following argument is valid.

The Lawyer Argument

If he has a good lawyer, then he will be acquitted.
 He has a good lawyer.
 Therefore, he will be acquitted.

But there seem to be some grounds for doubt whether this recommendation is right. Suppose both premises are true. Suppose the first premise means that if he has a good lawyer then, on balance, but subject to other considerations, he will be acquitted. Then does it follow that the conclusion has to be true? Well, not necessarily, it seems. You could have a good lawyer, but the other side could have an even better one. Of course, one can argue about the first premise, and what it should really be taken to mean. Some would say that you could interpret it in such a way that the hook adequately renders its meaning. But on the whole, this argument has never seemed all that plausible.

Of course, as in the Emil argument, and the other two arguments cited above, the CFL will argue that the major (conditional) premise cannot really be true if the conclusion is false. For example, suppose that in the case of the lawyer argument, the other side has a better lawyer who, as things turn out, wins the case. Does that event defeat the *modus ponens* type of argument above? Well no, according to the CFL, because the conditional premise of the argument, as things turned out, was not really true. The premise was not true under all circumstances, and therefore (according to the CFL), it was not "true" in a strict, objective sense. The assumption is that there is no sense of "If *A* then *B*" in English in which it is now or was ever true that he has a good lawyer and will not be acquitted. According to the CFL then, *modus ponens* arguments are always valid, and always have to be valid. The CFL will argue that you need to interpret the conditional in the major premise in such a way that it can only be (strictly speaking) true if it is true under all circumstances. This argument takes us to a consideration of the subject of conditionals. The CFL argues that conditionals, like those in the arguments above, can only be true if they are true under all circumstances, and otherwise they are false. Based on a discussion of conditionals, it is argued below that in all of the cases cited above, and similar cases, the conditional does not need to be treated exclusively in this strict way.

It is rare to doubt the validity of *modus ponens*. However McGee (1985) has

presented several counterexamples to *modus ponens* that have the form of the argument in the case below.

The Election Case

Opinion polls taken just before the 1980 election showed the Republican Ronald Reagan decisively ahead of the Democrat Jimmy Carter, with the other Republican in the race, John Anderson, a distant third. Those apprised of the poll results believed, with good reason:

If a Republican wins the election, then if it's not Reagan who wins
it will be Anderson.
A Republican will win the election.

Yet they did not have reason to believe:

If it's not Reagan who wins, it will be Anderson.

The form of the argument used in the election case is that of *modus ponens*, but it is a special instance in which one of the component propositions is itself a conditional. Is this a kind of odd case, in that other cases of *modus ponens* that do not have this special nested feature are valid? McGee seems to think so. He takes his counterexamples to show that "modus ponens is not an entirely reliable rule of inference." (p. 463). The presumption seems to be that, apart from such special cases, *modus ponens* is a deductively valid inference.

The doubts that will be expressed below about the validity of *modus ponens* are quite different from those of McGee. In fact, the highly specific doubts expressed by McGee show that in general he does accept *modus ponens* as a valid form of argument, without having the much more general kind of doubts that will be expressed in the arguments below. It will be shown below that there are many ordinary cases of arguments that have the form of *modus ponens* but that are deductively invalid. In recent work on argumentation schemes (Hastings, 1962; Perelman and Olbrechts-Tyteca, 1969; Kienpointner, 1992; Walton, 1996), many examples are presented that show that such arguments are not rare or unusual in any way. In fact they are extremely common. In (Walton, 1996) the thesis is promoted that many common arguments that have the *modus ponens* form (or what appears to be the *modus ponens* form) are presumptive arguments that can be judged to be structurally correct or incorrect by the standards appropriate for presumptive arguments. These arguments can be structurally correct, from a presumptive point of view, but they are not deductively valid. The forms of such arguments are called argumentation schemes. Many would class these types of argument as abductive, or as species of abductive arguments. The first premise of many of these arguments is a conditional. But it is not a "tight" conditional of the kind that is falsified or defeated by a finding of one instance where the antecedent is true and the consequent is false. It is a kind of loose conditional asserting that if the antecedent is true then generally the consequent is true, subject to exceptions. In current usage, it is a defeasible or default type of conditional.

3. Defeasible Inferences and Modus Ponens

Of course, there is a question here of how to define the *modus ponens* form of inference. As noted above, the typical argumentation scheme links a conditional warrant with another statement that fits the antecedent of the conditional. The conclusion then states the consequent of the conditional. But some might say that such a type of inference does not have the form of *modus ponens*. Why not? Well, they might argue that the first premise isn't really a conditional. Why not though? It has the form "If *A* then *B*." But the defenders of this view might counter that it is not a material conditional, of the kind appropriate for logic, because it is not the sort of conditional of the kind represented by the truth-functional connective called the hook (horseshoe). On the other side, many would say that there are lots of conditionals, like counterfactuals and probabilistic conditionals, for example, that do not have the form of the hook either. On this view, which is my own view of the matter, the if-then statement in the inference can rightly be classified as a species of conditional, even though it is not well represented by the truth-functional hook. On this view, the inference does have the *modus ponens* form. What needs to be appreciated is that whether an inference counts as having the *modus ponens* form depends on what counts as a conditional (that fits into the conditional premise). Thus in Section 3, I discuss some relevant aspects of the subject of conditionals.

There are some problems of usage in describing conditionals, and inferences based on them. The usage in describing these conditionals is not standardized between logic and computer science. In computer science, abductive conditionals of this sort, or sometimes even all conditionals used in inferences, are typically called "rules". This usage is somewhat comparable to Toulmin's terminology, in which a conditional of the sort cited above would be called a "warrant" (Toulmin, 1958). An example of how an artificial intelligence system uses types of inferences that are similar to *modus ponens* inferences is the ArguMed System constructed by Bart Verheij to provide argument assistance for lawyers. The ArguMed System uses what are called "step warrants" that express how particular statements can be used as a reason to support another statement. According to the account given by Verheij (1999, p. 44), a step warrant is similar to a Toulmin warrant, but also "plays a role analogous to the classic rule of inference *modus ponens*." How the system models an argument can be illustrated by the example of defeasible reasoning presented by Verheij (1999, p. 45):

Peter has violated a property right. As a result, at first sight, he has committed a tort. However, there is a ground of justification for Peter's act. As a result, on second thoughts, Peter's violation of a property right does not justify that he has committed a tort.

In the ArguMed System, as in many AI systems, conditionals are based on so-called rules. An example of a rule in ArguMed System is: "As a rule, if Peter has violated a property right, then he has committed a tort" (Verheij, 1999, p. 45). But

the key feature of rules, and of the kinds of inferences commonly based on them, is that they are defeasible. What this means is that such rules are subject to exceptions so if the case proves to be an exception the inference will be defeated, or fail. This property of defeasibility is illustrated in the case cited above. As a general rule of law, if Peter has violated a property right, then he has committed a tort. But in this particular case, there is a ground of justification for Peter's act. In this case, therefore, there is an exception to the rule, and the inference is defeated. It does not follow that Peter committed a tort, as it seemed at first.

This process of basing an AI system on the model of defeasible reasoning is very common in AI systems. In fact, it is quite typical of how rule-based systems work. But the question posed here is how the model of reasoning in such systems is based on *modus ponens*, or some form of inference that appears to be quite similar to *modus ponens*, but is not quite the same. The kind of reasoning used in ArguMed does seem to depend on a central inference type that is similar to *modus ponens*. For example, from the account above, it can be seen that the following kind of inference would be used in ArguMed.

If Peter has violated a property right, then Peter has committed a tort.
Peter has violated a property right.
Therefore Peter has committed a tort.

This inference looks to all the world like an instance of *modus ponens*. And it seems to function in somewhat the same way as an inference. If both premises are true, then that is reason for accepting the conclusion as true. Yet it is also different from *modus ponens*, or at least the kind of *modus ponens* we are so familiar with in deductive logic. For in that kind of *modus ponens*, it is impossible for the premises to be true and the conclusion false. In ArguMed, that feature does not apply to the *modus ponens*-like inference just above. In that inference, it is possible for the premises to be true and the conclusion false in certain cases. These are the cases that defeat (or "undercut" in Verheij's terms) the inference. In this respect, the inference above is different from *modus ponens*, and is not properly classified as fitting the *modus ponens* inference form. In this respect, the arguments in ArguMed and AI are similar to the kinds of arguments studied in the literature on argumentation schemes. They look like arguments having the *modus ponens* form, and yet they are clearly quite different from deductive arguments in how they work.

The question then is whether this common and important form of defeasible inference in AI systems should be included as being of the *modus ponens* form or not. Judging from the way it is described by Verheij in ArguMed, it looks like it is not being classified as having the form of *modus ponens*. A step warrant, in Verheij's terms, plays a role in ArguMed that is only analogous to the material implication relation in *modus ponens*. What seems to be suggested is that the inference about Peter above does not have the form of *modus ponens*. Instead it has a form that is distinctive in its own right as a type of inference, but that is similar to *modus ponens* in the way it works. It moves forward by detaching the consequent of the

conditional as the conclusion, provided the other premise, which is also the antecedent of the conclusion, is accepted as true.

The basic issue at dispute can be formulated simply, using the following example of an inference. This inference is ubiquitous in computer science.

The Tweety MP Inference

If Tweety is a bird, then Tweety flies.
Tweety is a bird.
Therefore, Tweety flies.

My contention has two parts. The first part is that the Tweety MP inference has the form of *modus ponens*. The second part is that the Tweety MP inference is not deductively valid. The first part of the contention is that the Tweety MP inference has the following form of argument, for propositions *A* and *B*.

Form of Modus Ponens

If *A* then *B*
A
Therefore *B*

The Tweety MP inference appears to have the form of *modus ponens* on the surface, but what is under the surface? Of course, it is quite possible that once we know more about how the Tweety MP inference works, we will see that it will be shown to have a structure that is not that of *modus ponens*. The Tweety MP inference could be represented as having a form with its own special premises. One premise states a general rule in abstraction from applications to a specific case, while other premises are based on applications to the specific case. In (Walton and Krabbe, 1995, p. 180), the following form of inference is put forward as a way of representing a nondeductive inference of the kind associated with default reasoning. Let's call this form of inference form*.

Form of Inference*

(1) (rule *n*): $B_x \Rightarrow F_x$
(2) B_t
(3) Rule *n* applies to the present case
Therefore F_t

Form* of inference has as its first premise a rule (rule *n*) stated at a level of general abstraction. It says that generally, if something has property *B* it will also have property *F*. The second premise states that individual *t* has property *B*. Normally, by deductive reasoning, the conclusion that individual *t* has property *F* could be detached. But rule *n* has only been stated at (1) in the abstract. It might fail to apply in this case. However, once premise (3) is in place, we know that rule *n* applies to the case of *t*. Hence the conclusion that *t* has property *F* can be detached. Form*

can be taken to represent the form of a default inference like that in the Tweety case, where *B* stands for the property of being a bird, and *F* stands for the property of being something that flies.

It seems evident that form* is not the same as the *modus ponens* form of inference. It is more complex, has more premises, and the conditional in the first premise is not the absolute conditional (or generalization) of deductive logic. If so, there is no problem about the universal validity of *modus ponens*. Although form* may not be a deductively valid kind of inference, that is no problem, because it is not a *modus ponens* form of inference. But then we must consider inferences like the Emil argument. Is it really a defeasible inference of form* type, or is it really a *modus ponens* form of inference? If the Emil inference, and the many inferences like it cited as having the *modus ponens* form in the logic textbooks, is not of the *modus ponens* form after all, then there is no big problem, except that many of the examples used in the logic textbooks, and how they are treated, will have to be changed. In the case of the Emil argument for example, the inference will need to be portrayed as not being of the *modus ponens* type, but as having a form expressed by form*, or something like it. For example, the first premise will need to be analyzed as saying something like, "If Emil is captured then (all else being equal, and barring unusual developments), all will be lost." But of course, such an analysis is controversial. We seem to be dealing with some sort of conditional in the Emil argument, but what sort? Is it a default conditional, a probabilistic conditional, or something more like the material conditional of deductive logic? A more general issue is the question of whether the proposition "If Tweety is a bird, then Tweety flies" is a conditional. Those who accept that view that all inferences having the form of *modus ponens* are valid may maintain that although this proposition looks like a conditional, really it is not. Alternatively, they might argue that although it is a sort of conditional, it is a special kind of conditional that is not really proper to have as a premise in a *modus ponens* type of argument.

Now we need to consider the second part of my contention, that the Tweety MP inference is not deductively valid. I contend that it is possible for the premises to be true and the conclusion false. Consider the case in which Tweety is a penguin. In this case, it is true that if Tweety is a bird then Tweety flies. But it is false that Tweety flies. The CFL will, following the TF principle, challenge the first claim. She will argue that if Tweety really is penguin, a type of bird that does not fly, then the conditional in the major premise of the Tweety argument must be false. The CFL will feel that, in such a case, there can be no conditional in which it was or is true that if Tweety is a bird then Tweety flies. In the CFL's view, Tweety can't fly, even though he is a bird. According to the CFL, that's the end of it. The conditional simply can't be true, and thus *modus ponens* cannot be invalid. But is that the end of it? No, not really. The debate about the validity of *modus ponens* has now been transformed into a debate about conditionals, and about the truth-conditions or acceptability-conditions for conditionals. The debate is now about whether only certain types of conditionals are appropriate to be used in genuine *modus*

ponens arguments. One standard line that will be very appealing to the CFL is the line that in a real *modus ponens* inference, the conditional premise must be a special kind of conditional. The defining characteristic of this type of conditional is that it is always false when the antecedent is true and the consequent is false. As applied to the Tweety MP inference, of course, this analysis will make the inference come out as deductively valid. But in my view this analysis is open to doubt. The reason I hold it to be open to doubt is that I maintain that the conditional premise in the Tweety MP inference can be true (or reasonably accepted as holding true), even if it is a fact that the antecedent is true and the consequent false. This view will seem peculiar, or flat out wrong, to the CFL. To get to the point where my view is even arguably plausible, we have to consider the nature of conditionals. Much has been written about conditionals, and the subject is convoluted and controversial (Sanford, 1989). In Section 4, a particular approach to conditionals is outlined. In this approach, it is shown how the subject of conditionals is linked to the subject of generalizations.

4. Conditionals and Generalizations

Let's begin with a generalization about how logic textbooks treat generalizations. Although they do write very briefly about generalizations in an informal manner, in treating fallacies like the fallacy of hasty generalization, they don't write very much about generalization *per se*. And when they do mention them, they tend to see them as universal generalizations, that make a claim about all members of a class, as represented by the universal quantifier, "For all x ." For example, the most widely used logic textbook (Hurley, 2000)³ only mentions generalization once (outside the informal fallacies chapter). A general statement is defined as a statement that "makes a claim about *all* the members of a class" (p. 38, Hurley's italics). The next most popular logic textbook, Copi and Cohen (1998), define 'generalization' in their Glossary and Index of Logical 'terms, as "In quantification theory, the process of forming a proposition from a propositional function, by placing a universal quantifier or an existential quantifier before it" (p. 691). But on the same page, the following definition for "generalization, inductive" is given: "The method of arriving at general or universal propositions from the particular facts of experience." What is suggested by the Hurley account is that there is only one kind of general statement, the kind I would normally call the universal generalization. But is this logical viewpoint consistent with practices of argumentation? What about my generalization about how logic textbooks treat generalizations? Is it a universal generalization, falsified by a single counter-example? I think not. So can there be different kinds of generalizations? The Copi and Cohen approach is more liberal, allowing for three different kinds of generalizations—two in quantification theory and one inductive type. Perhaps it is too liberal. An existential quantifier in deductive logic is made true by only a single instance being true. It comes out as true even if one thing that has one property also has another property. Is that kind of

statement a generalization? Not the way I would prefer to use the term. Copi and Cohen also define generalization as a process or method. This too is less than helpful. I would like to see a generalization defined as a type of statement or proposition.

According to a view I have argued for elsewhere (Walton, 1996. Chapter Five), there are three kinds of generalizations. Distinguishing these three types of generalizations is vital to getting a grasp of the fallacy of hasty generalization. But why is this view about generalizations relevant to the validity of *modus ponens* inferences? The answer is that how logic should treat conditionals is closely linked to how logic should treat generalizations. Corresponding to the three types of generalizations, there are three types of conditionals. This correspondence turns out to be quite important with respect to *modus ponens*, because it reveals that there are also three kinds of *modus ponens* inferences, depending on the type of conditional in the inference. Let's start with the three kinds of generalizations.

The first type is the (absolutely) universal generalization, of the form "All *F* are *G*," for classes *F* and *G*. The defining characteristic of this type of generalization is that it is falsified by one counter-example. If you can find one thing that has property *F* but does not have property *G*, then the universal generalization "All *F* are *G*" is false. For example, the generalization "All men are mortal" is falsified by a single instance of a man who is not mortal. We could call this *the single counter-example characteristic*. Both the other two types of generalizations fail to have this single counter-example characteristic. The second type of conditional is the probabilistic conditional of the form, "So many *F* are *G*," where the quantifier 'so many' is represented by a fraction between zero and one, representing the probability of something that has property *F* also having property *G*. An example is the generalization, ".72 of blue penguins live in New Zealand." The third type of generalization is the plausibilistic generalization, of the form "*F* are generally *G*, but subject to exceptions." An example is the generalization "Birds fly."

The absolutely universal generalization is familiar from deductive logic. For example in the A-proposition of syllogistic logic, or the universal quantifier. The universal generalization is expressed in modern deductive logic as a conditional of the form, $(\forall x)(Fx \supset Gx)$. This way of modeling the universal generalization exhibits the connection with the conditional. The universal generalization is defined in terms of the material conditional (hook). This way of defining it makes sure that it has the single counter-example characteristic. But it also indicates why the material conditional would not be suitable for defining either of the other two types of generalizations. For neither of these has the single counter-example characteristic. At present it is controversial whether the plausibilistic generalization can be reduced to some form of probabilistic generalization, or whether it is a distinctive type in its own right.

It was noted by Alfred Sidgwick how logic sometimes assumes that only the universal generalization can "properly serve as ground of inference" (1893, p. 23).

But, as Sidgwick also observed, "it is comparatively seldom in actual argument—never, perhaps, where a really difficult or disputed question is raised—that we are able to rest our case in a single faultless generalization, like 'All men are mortal'" (*ibid.*). Sidgwick was a precursor both of informal logic and of the recent flowering of artificial intelligence research on defeasible reasoning. He recognized the importance of defeasible generalizations, and the inferences based on them, in practical reasoning in everyday argumentation. But Sidgwick was a voice in the wilderness. The dominant practitioners of logic, both during and after his lifetime, stressed the universal generalization and paid little or no attention to the plausibilistic kind of generalization that is open to exceptions. However, opinion is now shifting somewhat more towards Sidgwick's view. Josephson and Josephson wrote: "The universal quantifier of logic is not the universal quantifier of ordinary life, or even of ordinary scientific thought" (1994, p. 30). In their view, "reasonable generalizations are hedged" (subject to exceptions). Moreover, they added: "Predictions from hedged generalizations are not deductions" (p. 23). This view is well supported by the recent computational literature on defeasible reasoning in artificial intelligence.

It is my view that there are three types of conditionals, just as there are three types of generalizations. The first type, the absolute conditional, has the characteristic that if the antecedent is true and the consequent is false, the conditional is false. The absolute conditional transfers truth from the antecedent proposition to the consequent proposition. If the antecedent is true, then the consequent is true, and there is no doubt about it. The material conditional (the hook of classical logic) is a subspecies of absolute conditional. The second type of conditional is the probabilistic conditional. Instead of being defined in terms of truth values, it is defined in terms of probability values. It transfers probability from the antecedent to the consequent, but leaves doubt open as to whether the consequent is true (even if the antecedent is). If the antecedent has a certain degree of probability then probably the consequent also has a certain degree of probability (depending on the probability of the antecedent and the probability of the conditional link between it and the consequent). The third type of conditional could be called the abductive, defeasible, or plausibilistic conditional. It is defined in terms of plausibility values. It throws a weight of plausibility from the antecedent to the consequent. A weight of plausibility in favor of a proposition's being acceptable only gives a tentative reason for accepting that proposition, subject to doubt, and subject to potential retraction. A weight is generally a small amount of evidence that is not very significant in itself, but that, taken together with a larger body of evidence, tilts the burden of proof to one side or the other on a balance of considerations. An example is the conditional, "If Tweety is a bird, then Tweety flies." If the antecedent is accepted in a given case, then a weight of plausibility is shifted to the consequent's being rationally accepted (other things being equal). But it must be emphasized that this kind of conditional is inherently subject to default. If new information comes into the body of evidence in the case indicating strongly that

Tweety is penguin, the conditional no longer gives a reason for accepting the consequent, even if the antecedent is accepted.

For example, the conditional, "If Socrates is a good man, then Socrates is a man," is an absolute conditional. It can be classified this way because it is not possible for the antecedent to be true and the consequent false. Thus there is no need to add qualifications to this conditional in order to accept it as true. Basing an inference on this conditional as warrant can yield a *modus ponens* type of inference that is deductively valid. The following argument has the *modus ponens* form and is deductively valid.

The Socrates MP Inference

If Socrates is a good man, Socrates is a man.
Socrates is a good man.
Therefore, Socrates is a man.

If you examine the Socrates MP Inference, you can see that there are no exceptions that might defeat the conditional premise. The conditional is absolutely acceptable without reservations. Contrast this inference with the Tweety MP Inference. It is rightly classed as a defeasible inference. The reason for this classification stems from the conditional that is the warrant of the inference. It is defeasible, and hence the inference itself is defeasible. But the Socrates MP Inference is not defeasible. It is deductively valid. It is not possible for both premises to be true and the conclusion false. This is a "tight" or exceptionless kind of inference, one rightly associated with deductive reasoning. How an inference should be classified thus depends on the generalization or conditional that functions as the warrant of the inference.

5. Abductive Inference and *Modus Ponens*

The issue of different types of generalizations and conditionals has now raised a closely related but highly controversial issue. This issue is the logical form of abductive inferences. Abductive inference is often taken to be the same as what is called inference to the best explanation. Both of these types of inference are strongly associated with a type of argumentation called argument from sign. An argumentation scheme for argument from sign is presented in (Walton, 1996, pp. 47-49). These three different categories are frequently run together and confused, one of them being taken to be equivalent to one or more of the other two. Because abductive inference is crucial to recent developments in artificial intelligence, there is a huge (but confusing) literature on the subject in computer science. These three categories have already been extensively discussed and analyzed in (Walton, 1996, Chapter 8). Rather than cover all that territory again, the reader can be referred to this chapter. But what can be done here is to try to sketch out how the issue of *modus ponens* always being deductively valid relates to all these other issues.

It is best to begin with a simple example of a very common kind of inference.

The Measles Inference

If a patient has red spots (of a certain kind), then the patient has measles.

This patient has red spots (of this certain kind).

Therefore, this patient has measles.

The measles inference appears to have the form of *modus ponens*, on the surface. Perhaps under the surface, it may really have a more complex form like form*, that may be a special subtype of *modus ponens* inference, or may not. It is a typical kind of inference very commonly used in medical diagnostics (Fox and Das, 2000). It can also be classified as an instance of argument from sign. A sign or indicator is observed, and then since the sign is linked to a certain condition, or as identifying something, that condition or thing is concluded to be present.⁴ The first thing that is immediately obvious about this inference, and this kind of inference generally, is that it is not deductively valid. It is possible for the premises to be true and the conclusion false. It may turn out to be the case that when tests are done, the tests show that the patient does not have measles. So what good is the inference if the conclusion turned out to be wrong? The function of the inference is to make a guess or hypothesis that can lead to testing. Once the tests are in, the findings may confirm the guess, or they may show it was false. Either way, knowledge is gained about the patient's diagnosis. If the initial guess can be ruled out, then other diagnoses can be explored and tested. If the guess turned out to be right, then treatment for measles can be undertaken, and the possibility of having to deal with other possible diseases can be set aside. So even though the inference is not deductively valid, it performs a very useful function as a kind of reasoning in medical diagnosis.

Let's take a closer look at the conditional in the first premise. What it means is that if the patient has these kind of red spots, then it is a good hypothesis that the patient has measles. The presumptions are that no other contra-indicating evidence is available in the case at this point in the diagnosis, and that these are the right kind of red spots that look like the kind that normally indicate measles. What does the 'if-then' mean then? It doesn't mean that if the antecedent is true, the consequent must be true. It just means that if the antecedent is true, then everything else being equal at this point in the investigation of the case, the consequent is a good working hypothesis to go ahead with, at least as a basis for conducting tests, or if tests are not necessary, as a basis for provisional action or inaction. The conditional has the function that if the antecedent is triggered, a line of actions and further investigations is laid out by moving to the consequent as a working hypothesis. The measles inference is better construed as having the function of narrowing down the search, rather than putting an end to it. Thus conceived, while the *modus ponens* form of inference exemplified in the measles inference

seems to be quite reasonable, and to have a legitimate and useful function as a kind of reasoning, it is not conclusive. It is not a logically necessary kind of inference of the kind associated with deductive logic. Indeed, to appreciate how it works correctly as a useful inference, you need to appreciate that it is not deductively valid. To see it as deductively valid would be a kind of fallacy—perhaps a fallacy of dogmatic thinking that fails to be sensitive to exceptions to a rule. In traditional terms, it could be classified as a kind of hasty generalization fallacy, or *secundum quid* fallacy of overlooking exceptions to a rule. It is most likely that the measles inference would nowadays be classified in AI as an abductive type of inference.

Peirce (1965, p. 375) defined abduction as a process of forming a hypothesis that can be used as a tentative step in an investigation to explain some observed data. He saw abduction as a process of inference to the best explanation, "where we find some very curious circumstance, which would be explained by the supposition that it was a case of a certain general rule, and thereupon adopt that supposition." Peirce gave the following example of a finding that could be the basis of an abductive inference: "Fossils are found; say, remains like those of fishes, but far in the interior of the country" (*ibid*). This finding poses a question. How could fish be present where there is no water? To explain the finding, and thereby answer the question, a hypothesis is postulated that the sea once washed over this land. The conditional part of such an inference alluded to by Peirce is the use of a "general rule." What is this general rule? Presumably, it is a rule to the effect that where you find fish, you can expect to find water. The generalization on which the rule is based could be that fish normally live in water. So unless there is some other explanation that accounts for it, if you find fish, then you should find water in that area.

What would now be called abductive inference was not only shown to be vitally important in science by Peirce. How inference to the best explanation is a crucial kind for understanding how logical reasoning as used in the law of evidence was shown by Wigmore (1940). Above it has also been indicated how abductive reasoning is extremely common in medical diagnostics. Since abductive reasoning is not deductively valid, it has become crucial to try to see what form it has as a type of argumentation. What is clear from the above examples and accounts is that it is some type of inference to the best explanation. What is meant by this phrase is not inference to the best possible (or maximal) explanation. What is meant is that several explanations of a finding or phenomenon could be possible. The abductive inference narrows down the range of possibilities by weeding out the less plausible ones, narrowing down the search to more plausible ones, or even perhaps to one that stands out as highly plausible. But it does not rule out all the other possibilities except the chosen one. It is not deductively valid, and as indicated above, it would be a kind of fallacy to portray it in this way.

The following account of the form of the abductive type of inference is given by Josephson and Josephson (1994, p. 14):

(*Abduc.Inf*) : *D* is a collection of data.
 Hypothesis *H* explains *D*.
 No other hypothesis explains *D* as well as *H*.
 Therefore *H* is plausibly true (acceptable).

What is important to see about this form of inference, so represented, is that it proceeds by excluding other hypotheses. It cuts down the range of competitors. The form given by Josephson and Josephson even indicates that the one selected hypothesis—the so-called "best" explanation—rules out all the others. But this terminology is a bit misleading. What this form of inference does is to suggest a hypothesis that is plausible, because it stands out as an explanation that fits the case. But there could be other explanations that are equally plausible in some cases. In such a case, further tests may need to be devised that distinguish between the two explanations and picks one as preferable to the other. Or in some cases, two may be equally good explanations and you just need to pick one. In any event, the inference to the best explanation is at its most useful in the exploratory stages of an investigation, where it is useful to narrow down the range of hypotheses for testing by going with what seem (tentatively) the most plausible one or ones to pursue.

The abductive inference has two kinds of premises. One is a factual or observational kind of premise, where some observed data or object is classified as fitting a certain pattern. For example, red spots may be observed, and then classified as "red spots." This finding then functions in the other premise as one part of a conditional or generalization that links the finding to a hypothesis. When the two premises are put together in an abductive inference, the finding from the first premise fits into the second premise as antecedent, triggering the detachment of the consequent. The statement in the consequent is then drawn as the conclusion of the inference.

The subject of the form of abductive inference is highly controversial at the moment. Some would say that the conditional premise should be turned around. For example, in the measles inference above, the conditional premise should be turned around, so that the inference should be seen as follows.

If the patient has measles, then the patient will have red spots (of a certain kind).
 This patient has red spots (of this kind).
 Therefore, this patient has measles.

When the inference is seen this way, it may seem more like it could be stronger, perhaps even deductively valid. For the conditional could be seen as expressing a sufficient condition relation. But the inference has the form of affirming the consequent, a form of inference that is deductively invalid. Some think that abductive inferences characteristically have the form of affirming the consequent. The controversy over whether abductive inference is better seen as having the form of *modus ponens* or the form of affirming the consequent has been extensively dis-

cussed in (Walton, 1996a, Chapter 8). Rather than going over it all again here, it is better to refer the reader to that discussion of the issue. The issue is tied in with a fallacy that Aristotle called the "fallacy of consequent," which has to do with turning a conditional around and getting it backwards when using argumentation from sign. This issue has not yet been fully resolved, and is highly controversial. This is not the place to try to settle it, or consider all sides of it. Suffice it to say the viewpoint advocated here is that abductive inference is best seen as having a *modus ponens* form of inference that is deductively invalid.

The real problem here is that there is so much of an antipathy, from a viewpoint of traditional formal logic, to seeing inferences having the form of *modus ponens* as being other than deductively valid. This viewpoint seems to suggest that if abductive inferences to the best explanation can be seen as having some form other than *modus ponens*, some form that is not deductively valid, like affirming the consequent, then that must be the best way to model them. But a consideration of cases suggests otherwise. The considerations adduced above indicate that the conditional in an abductive inference goes from antecedent finding of fact (observation, indicator, sign) to a consequent that postulated a best explanation of the observed fact. This way of viewing the conditional premise yields an account of abductive inference that makes it have the form of *modus ponens*. But this analysis of abductive inference appears so unpleasant to the practitioners of formal logic that they do not want to consider it unless pushed to the wall. The current dogma that all inferences having the form of *modus ponens* are deductively valid is so powerful that the above analysis of abductive inference is not considered to be open to discussion. It is blocked out of consideration by the reigning dogma.

6. The TF Principle and the Prevalence of the Dogma

The reasons for the prevalence of this dogma are not hard to find. If students can be convinced that all arguments of the *modus ponens* form are deductively valid, they are impressed with the power of formal logic. Formal logic seems to apply to all kinds of common arguments used in everyday reasoning. Moreover, all kinds of highly problematic and potentially complicated questions about different kinds of conditionals can be avoided. But a little reflection on these everyday arguments can easily raise significant questions about the dogma.

When we assert a conditional of the form "If *A* then *B*," we mean that the truth of statement *A* is sufficient for the truth of statement *B*. But do we mean "sufficient" in the sense that if *A* is true then *B* must necessarily be true? In other words, do we mean that it is impossible for *A* to be true and *B* false? We do sometimes, but only in cases where the connection is tautological, as in the statement, "If this book is a red object then this book is a colored object." In many common kinds of cases, what is meant is that the truth of *A* is sufficient for the truth of *B*, all else being equal. This kind of sufficiency applies to presumptive conditionals like, "If Tweety is a bird, then Tweety flies." In such a case, the acceptability of the

antecedent is sufficient for that of the consequent, but only in an abstract or general way that is subject to exceptions in specific cases. If we know that Tweety is a bird, then we can accept that Tweety flies, unless we know that Tweety is a penguin or an ostrich (or otherwise that Tweety is an exception to the general rule). In still other cases, the notion of sufficiency is probabilistic. For example, if I say, "If I drop this piece of chalk, it will break," what I (most plausibly) mean is my dropping it will probably be sufficient for its breaking, as far as I can predict. When I make such an assertion, I am not saying that it is impossible both that I should drop the chalk and that it does not break. I know that this unlikely outcome is possible. But still I say, "If I drop the chalk, it will break." What is shown, so far, is that there are three notions of sufficiency, and each of them corresponds to a type of conditional. But now another factor needs to be taken into account.

What has been introduced into the way of thinking typical of formal logic is a fourth kind of sufficiency that can be called material or factual sufficiency. One statement *A* is materially sufficient for another statement *B*, in this sense, where it is not the case that *A* is true and *B* is false. There is a conditional corresponding to this notion of sufficiency. In section 4 above it was called the absolute conditional. Characteristic of the absolute conditional is what can be called the TF principle: if the antecedent is true and the consequent is false, then the whole conditional is false. The TF principle by itself is enough to make *modus ponens* deductively valid. The proof is as follows. Consider an argument having the form of *modus ponens*, where both premises are true.

If *A* then *B*
A
 Therefore *B*

Now we reason by the following deduction:

1. The second premise is true. So *A* is true.
2. But then consider the first premise. It is a conditional. So by the TF principle, if *A* is true and *B* is false, the whole first premise is false.
3. But the whole first premise, by supposition, is not false.
4. Hence it is not the case both that *A* is true and the *B* false.
5. But *A* is true, again by supposition.
6. Hence it is not the case that the *B* is false.
7. Hence *B* is true.

This proof shows that if the TF principle applies, then *modus ponens* must be a deductively valid form of argument. What is shown is that if we accept the TF principle then, as long as we continue to accept it, *modus ponens* must always be a deductively valid form of argument. And since the TF principle seems quite minimal, and reasonable, then it seems that *modus ponens* must be accepted as valid. So it seems to follow that if any argument has the form of *modus ponens*, it must be deductively valid. All this conventional wisdom of logic seems very reasonable. But we know it has to be wrong, if the theory above about the three kinds of generalizations and three kinds of conditionals is right. Where is the fly in the ointment?

The fly in the ointment becomes apparent once we start to examine the chain of reasoning in ordinary examples of *modus ponens* inferences. Consider the following argument that has the form of *modus ponens*:

If I drop the chalk, it will break.
I drop the chalk.
Therefore, the chalk will break.

Suppose the premises are true. Is it possible for both premises to be true and the conclusion false? That depends on the supposed truth conditions for the conditional premise. What is required to make this conditional come out true or false? The answer is that it depends on what you mean by it. What we do in the practice of teaching logic is to bring the TF principle to bear on the inference. We remind our logic students that the conditional is false if the antecedent is true and the consequent false. It follows from this interpretation (as spelled out above) that the *modus ponens* inference is deductively valid. But is that really what I mean when I say, "If I drop the chalk, it will break"? The problem is that what I say is ambiguous. It could mean that. But we could also plausibly take it to mean something else. We could plausibly take it to mean that if I drop the chalk, it will probably break. Even if it fails to break, my assertion that if I drop it, it will probably break, could still be true. But the interpretation of the conditional as a material conditional is possible as well. This interpretation can be supported by using the following line of argument. Suppose we frame the argument in the form of a bet. I bet that if I drop the chalk it will break. You bet that if I drop the chalk, it will not break. I drop the chalk. It does not break. So you win the bet. What I said turned out to be false. This line of argument supports the interpretation of the conditional as a material conditional to which the TF principle applies. But either interpretation is possible. It could be that when I claimed that if I drop the chalk it will (probably) break, I meant to allow for its not breaking as being quite compatible with my claim. Here the TF principle does not apply. But the usual way we teach deductive logic, we take the conditional to have the deductive interpretation by invoking the TF principle. "Then we can go ahead and treat all arguments of the form *modus ponens* as being deductively valid. This process seems reasonable, but in reality it is a logical leap based on an ambiguity.

The plausibility of the TF principle presumes a timeless and exceptionless use of the term 'true' that blocks out any argument against the validity of *modus ponens*. As noted above, the CFL takes the view that if Tweety is a penguin, then there is no sense in which it now or ever was true that if Tweety is a bird then Tweety flies. The same invocation of the semantic terms 'true' and 'false' blocks out the possibility that *modus ponens* could be invalid in any of the other cases. Take the lawyer argument for example. Suppose the lawyer for the other side wins the case. Does that mean the conditional in the lawyer argument is false, or at any rate not true? Well, as things turned out the conditional was defeated in this case. But in general, the conditional is still tenable. Yet for the CFL, it is inappropriate to say it is "true", in her timeless sense of the term. Breaking away from this way of

describing arguments however, it could be said that the conditional is acceptable, as a general rule, but it defaulted in this case. What this discussion shows is that when evaluating defeasible arguments, it may be better to shift from the timeless language of truth to the pragmatic language of acceptance, or commitment as it is often called (Walton and Krabbe, 1995). The lawyer argument does present premises that give a good reason to support the conclusion. But it can be misleading to say that the premises can only be evaluated as true (or false) in some timeless sense of these terms. Better to evaluate the argument by saying that if the premises are acceptable, then the conclusion should be acceptable too, all else being equal. Invoking a strict meaning of the term 'true' seems to remove the problem. It does so by assuming that the lawyer argument, and the other arguments cited above, could never really be arguments of the *modus ponens* form, because where they do fail, the major premise could never really be said to be true. But invoking such a strict sense of the term 'true' is a way of refusing to see the argument in any other than deductive terms. It is to insist that if the antecedent actually turns out to be true and the consequent false, the conditional could never really be "true" in the absolute sense anyway. Thus the whole idea that *modus ponens* could ever be invalid can simply never arise. On the CFL's view, you could never doubt the validity of *modus ponens* using the kinds of counter-examples set out above. But the CFL's view doesn't deal with the reality of such arguments.

The problem with the CFL view is that despite its pedagogical attractiveness, and its logical provability (as shown above), it conceals an important ambiguity. It interprets the conditional as expressing a material sufficiency condition. Or at any rate, it only allows a set of adequacy (truth) conditions for conditionals that can be expressed by the holding to the TF principle. The TF principle guarantees the validity of *modus ponens*. But is that interpretation what we always mean when we assert a conditional? No, it is not. As shown above, in many cases the conditional is only meant to express a probabilistic sufficiency condition. Many conditionals only assert that if the antecedent is true the consequent is probably true. The typical conditional really says that if the antecedent is true in a given situation, and all other factors are held constant in that situation, then the consequent is also true (or will be). These truth conditions for the truth of the conditional are compatible with the antecedent being true and the consequent false. But probabilistic conditionals are not the only kind to be considered. There are also the presumptive (abductive) kinds of conditionals to be considered.

In a presumptive conditional, the conditional can still hold, even if the antecedent is true and the consequent is false in a given instance. It can happen where something fails to hold constant. The "all other factors being held constant" factor in the conditional means that the truth or falsity of the conditional is not a function of the truth values of its components. Its truth or falsity also depends on these other factors that are presumably held constant. The material interpretation of the conditional in a *modus ponens* argument overlooks this default way of interpreting the conditional. That can appear quite reasonable, for, after all, the truth-functional

interpretation is simpler and absolute. It satisfies our longing for precision and simplicity, all in one attractive package. But our selection of it is really based on an ambiguity. We take the simpler choice and pass by the more complex and messy interpretation of what asserting the conditional commits the assertor to defending. That choice is not unreasonable, as long as we are aware that it is a choice. But we are now so used to the deductive interpretation, and it has become such an entrenched dogma, that many of us are no longer aware that there is an ambiguity, and that a choice is there to be made.

We now come back to the subject of why *modus ponens* inferences are not deductively valid in presumptive inferences. The case in point is the Tweety MP inference, displayed again below.

The Tweety MP Inference

If Tweety is a bird, then Tweety flies.
Tweety is a bird.
Therefore, Tweety flies.

In evaluating this inference, we first of all need to ask: under what conditions should it be evaluated as structurally correct or incorrect? One answer proceeds by applying the TF principle, making the inference deductively valid. On this interpretation, since the argument has the form of *modus ponens*, it is impossible for the premises to be true and the conclusion false. But as the analysis of conditionals in the previous section has shown, this condition is not appropriate for evaluating the Tweety MP Inference. It is in fact possible for both the premises to be true and the conclusion false. The reason is that it could turn out that Tweety is a bird that does not fly, like a penguin. This possibility could be made real if, in fact, as the case proceeds, it turns out that Tweety is in fact a penguin, even though we didn't know it before. The first premise in the Tweety MP inference has the form of a plausibilistic conditional. Such a conditional leaves open the possibility that Tweety might turn out to be a non-flier. So by its nature, the Tweety MP inference is defeasible. It gives a reason for accepting the conclusion, based on acceptance of the premises, thereby throwing a weight of plausibility from the premises to the conclusion. But that acceptance is only tentative, subject to potential retraction if new information comes in that cancels out the old conclusion and leads to acceptance of a different one.

This way of evaluating the Tweety MP inference seems quite straightforward to those who familiarly deal with defeasible argumentation of the kind so commonly used in fields like computing and law (Prakken, 1997; Verheij, 1999). But many who most familiarly deal with deductive formal logic can be expected to resist it quite strongly. They will analyze the Tweety MP inference as follows. What you have to ask is whether the conclusion is true if the premises are true. Suppose that the premises are true. But what does it really mean to say that the first premise is true? For that premise to be true, it must be true that if Tweety is a

bird then Tweety flies. But what does that mean? It means that if "Tweety is a bird" is true, and "Tweety flies" is false, then the whole conditional must be false. So, in other words, if this premise really is true, and the other premise is also true, then the conclusion of the Tweety MP inference must be true too. This analysis represents the way defenders of the deductive validity of *modus ponens* really feel about this kind of inference. If the premises really were true (even contrary to fact), then the conclusion would have to be true too. By invoking the TF principle, they use the standard of the material conditional to evaluate all conditionals, thereby guaranteeing the universal validity of *modus ponens*. But as shown above, once this dogma is shown to rest on an ambiguity, its hold on us is broken.

7. Implications for Logic of Challenging the Dogma

My reasons for questioning such a central and widely held dogma of logic are not the negative motivation of revealing an error in the teaching of logic. The dogma is quite natural at a certain stage in the development of logic, and the reasons for holding to it are quite understandable, as explained above. My reason for critically questioning the dogma is to try to loosen up formal logic a little, so the subject can extend to the study of argumentation schemes that do not conform to the deductive model. Deductive logic deals with absolute conditionals and absolute generalizations. It represents a simple and precise but absolutistic way of thinking. There is nothing wrong with that. But much of the thinking used in everyday conversational argumentation is not absolutistic. It requires flexibility, and openness to retraction of commitments as new information comes into a case (nonmonotonic reasoning). A logic that is applicable to thinking in real cases must deal with conditionals that are not absolute. Although the applied part of logic is often categorized or dismissed as "informal," new developments in this area (Kienpointer, 1992; Walton, 1996; Reed and Walton, 2001) suggest that many of the most common kinds of arguments do have distinctive logical forms. These forms of argument, or argumentation schemes as they are now often called, do have a definite logical structure, even though they represent a kind of reasoning that is not absolutistic in nature. But it is hard to get logic to loosen up a bit, to get beyond absolutistic reasoning as the exclusive domain to be taken seriously, and to do research on argumentation schemes as non-absolutistic forms of argument that deserve serious study. One important reason is that, clinging to the dogma exposed above, logicians see them as all having the form of *modus ponens*. Hence any further study of the properties of their individual logical forms seems unnecessary. Ultimately, I believe that once the dogma is exposed, the way will be opened for a new branch of applied logic to study argumentation schemes for presumptive reasoning. The current problem is that doing work on common forms of argumentation within informal logic very often turns out to be based on systematically violating the dogma described above.

Recent work on argumentation schemes often seems to come very close to assuming the thesis that not all inferences having the form of *modus ponens* are deductively valid, and seems to move forward on that assumption. But to make the assumption explicit, by stating it as a general principle, is not something that could be easily undertaken. It just seems like too difficult a thesis to argue for, given the contrary view so often expressed or assumed in the field of logic generally and in logic textbooks. But sooner or later, as work on argumentation schemes in informal logic starts to converge with work on defensible argumentation in computer science, the assumption will become more and more explicit. The need to confront the issue was raised most explicitly by a recent review (Blair, 1999) of some work on argumentation schemes (Walton, 1996). What was questioned in the review was the apparent assumption that many of the abductive (presumptive) forms of argument (argumentation schemes) presented in (Walton, 1996) had the form of *modus ponens*, even though these schemes were not treated as being deductively valid forms of argument. But how could this be so? After all, the conventional wisdom in logic is that if these arguments have the *modus ponens* form, they must be deductively valid. As Blair (1999, p. 341) commented:

(S)everal of the formulations of argumentation schemes in Chapter Three represent valid argument forms, whereas Walton is quite explicit throughout the book that presumptive arguments are not deductive entailments. A case in point is the 'argument from popularity' or popular opinion which, Walton says, 'has the following argumentation schemes' (p. 33):

If a large majority (everyone, nearly everyone, etc.) accept *A* as true, then there exists a (defeasible) presumption in favor of *A*.

A large majority accept *A* as true.

Therefore, there exists a presumption in favor of *A*.

Besides the fact that 'defensible presumption' is redundant, this scheme has the form of *modus ponens*. Yet Walton says in the very next line: 'This kind of argumentation is deductively invalid...!'

These instances of associating argumentation schemes with the *modus ponens* form of inference were questioned by Blair as amounting to an apparent contradiction. The apparent contradiction noted by Blair turned on the assumption that argumentation schemes, like the one cited above, are not deductively valid, even though they have the form of *modus ponens*. After all, who could deny that all arguments that have the form of *modus ponens* are deductively valid? Blair assumed, with considerable justification, that all readers of the book would also accept this received view without questioning it. This assumption is evidence of not only how widely this view is accepted, but how it is even held as a dogma so that appearing to contravene it prompts a charge of contradiction.

As long as this dogma is in place, it appears that little or no room is left for the systematic investigation of argumentation schemes as a serious branch of logic. While many of these schemes do have *modus ponens* forms, the way the arguments actually function, it is clear that they are not deductively valid. Given the

current dogma, this apparent impasse leaves no room for further investigation of the kinds of inferential structure these arguments really have. One possible solution to the problem comes through the recognition that there are many common forms of inference that look like they have the *modus ponens* form, but that really, on closer inspection have a defeasible form of argument like form*. This solution, however, just seems to raise a different problem. How do you determine, in a given case, which form of argument is applicable, *modus ponens* or form*? The problem is that *modus ponens* seems to be the more general form. Form* is a special instance, it would seem. But arguments that have the form of form* are not deductively valid. The absurd result that seems to follow is that there are instances of *modus ponens* that are not deductively valid. This seems to be the same problem that has been plaguing the inquiry all along.

A way of salvaging this solution is to draw another distinction between two kinds of inference. In discussing a related problem, Verheij (2000, p. 5) has drawn a proof-theoretic distinction between two rule-based forms of inference.

Modus Ponens

Premises:

As a rule, if P then Q
 P

Conclusion:

Modus Non Excipiens

Premises:

As a rule, if P then Q
 P

It is not the case that there is an exception to the rule that if P then Q

Conclusion:

Q

What is most important to notice about these two rules is that they can be applied to two kinds of cases, cases where there are strict rules and cases where there are rules with exceptions. Accordingly, Verheij (2000, p. 5) adopts a policy that suggests the following general principle of applying the two rules to cases. In a case where both strict rules and rules admitting of exceptions might possibly come into play, *modus non excipiens* must always be used. In a case in which only strict rules are involved, *modus ponens* suffices as the appropriate rule of inference.

This policy seems basically right to me, from a point of view of applied logic, although I would use a slightly different terminology. I would say that there are two forms of argument involved in two different kinds of cases. Both forms of argument can be applied to cases in which the inference is based on a conditional warrant in the form of a general rule or generalization. Where the conditional is

purely strict, where no rules that admit of exceptions are involved, *modus ponens* in its deductively valid form can be used. For example in a theorem-proving case in geometry in which only deductive reasoning is involved (with universal generalizations about all triangles and so forth), the possibility that any of the general statements may be open to exceptions may simply not arise. In the other kind of case, the inference is based on a defeasible rule (or generalization) that admits of exceptions. Form* should be used for modeling defensible inferences of this kind. Verheij calls this form of argument *modus non excipiens*, but to avoid this expression, which is a little hard to pronounce, I would call it *defeasible modus ponens* (DMP). DMP should only be applied to certain special cases, those where the argument really has the structure of form*. DMP does not need be used for modeling strict inferences of the kind based on a rule (or generalization) that does not admit of exceptions. Within the confines of a case of this sort, there is no need to use DMP. *Modus ponens* (MP) will do. But considering cases of realistic argumentation in natural language discourse, DMP is clearly the model of choice.

Of course, some of the dispute is about terminology. The CFL will argue that DMP is simply not a form of *modus ponens* at all, and that therefore there is no problem. If an argument really has the form of *modus ponens* (MP), the CFL will argue, it is deductively valid, and that's the end of the issue. But I would reply that DMP is really a form of *modus ponens*, and that the term *modus ponens* should apply to its form. The reason is that DMP is a special form of MP, and DMP does have what should properly be seen as a *modus ponens* structure. Of course, the CFL might also argue that the major premise of the DMP type of argument, containing the defeasible generalization, cannot be properly said to be true (or false). So this form of argument does not have the MP form at all, and can never properly be said to be deductively invalid (or valid). The issue posed by this argument is partly terminological, and eventually it will be resolved by an emerging standard, as the practical urgency of the problem becomes more apparent to those teaching and conducting research in logic. The problem is that the current dogma does not do justice to the proper modeling of real argumentation.

Josephson and Josephson, as noted in Section 3, have also challenged the dogma that inferences based on generalizations of the kind used in everyday reasoning can be all analyzed using the absolute universal quantifier. Their arguments also raise doubts about the capability of deductive formal logic to model such inferences. As noted above, they doubt whether such inferences can be analyzed using the universal quantifier, as we are taught in logic class (Josephson and Josephson, 1994, p. 23). They do not see inferences based on the hedged quantifier of ordinary life as deductively valid. It could be presumed that they would also be dubious about modeling inferences based on conditionals that are "hedged" or open to exceptions, as having a deductive *modus ponens* form. Such a policy leads to an impasse. The only way out of this impasse is to challenge the dogma that all arguments having the *modus ponens* form are deductively valid. Once the dogma

is dissipated, the effect will be to open the way for an applied logic that can do justice to the analysis and evaluation of many common and important kinds of presumptive arguments. This step, in turn, will open the way towards a systematic and useful way of analyzing informal fallacies. It will also lead to a new union between informal logic and artificial intelligence, in which there is a matching between argumentation and the kind of dynamic, hedged, knowledge-based reasoning used in legal reasoning and medical diagnosis. In other words, challenging the current dogma opens the way to a new adjustment between what is now called formal logic and the field of informal logic (argumentation theory) that takes the study of fallacies, abductive inference and presumptive argumentation schemes as its subject matter. The first steps in this positive direction must surely arise from questioning the current dogmatic way of treating hedged inferences like those cited above.

The new view will restrict the applicability of deductive logic to *modusponens* arguments in which the conditional is of an absolutistic sort only. Deductive logic, on the new view, continues to be applicable to non-defeasible arguments of the *modus ponens* form, like the Socrates MP Inference. But it is no longer applicable to the evaluation of arguments that, like the Tweety MP inference, have a defeasible conditional as the major premise. These arguments require a different kind of reasoning structure, one that has its own standards for structural correctness of arguments, different from those of deductive logic. Deductive logic will remain a very important branch of logic—the first to be rigorously investigated, and in many ways, the most central part of logic. But informal logic will turn out to be an important branch of logic in its own right, a powerful argumentation engine that can be usefully applied to many of the most significant concerns of practical reasoning in computing and other fields like law and medicine. One reason lies in an empirical hypothesis that has not yet been tested, but that seems to be borne out by common experience. The hypothesis is that most arguments with the *modus ponens* general form encountered in everyday argumentation, and in fields like law and medicine, have the same defeasible structure as the Tweety MP Inference. The hypothesis implies that fewer of them have the same structure as the deductively valid Socrates MP Inference. If this hypothesis is correct, then it is informal logic rather than formal deductive logic that is most often applicable to *modus ponens* arguments.

Notes

¹ Of course, like all such disputes, much depends on how you define a key term. In this instance it is a question of how to define *modus ponens* as a form of argument.

² In using the expression 'argumentation schemes' here, we are primarily referring to forms of argument that are neither deductive nor inductive in structure. But this expression is defined more carefully below, and explained through examples.

³ On the back of Hurley (2000) it says that this book "is used by more students and instructors than any other throughout North America." But, of course, citing Hurley is an appeal to authority that could be wrong. The fact that this claim appears on the back of the textbook could also be classified as an *ad populum* appeal directed to potential buyers or adopters of the book.

⁴ See (Walton, 1996, pp. 46-47) for the argumentation scheme for argument from sign.

References

- Blair, J. Anthony. 1999. "Review of (Walton, 1996)," *Argumentation*, 13, 338-343.
- Copi, Irving M and Carl Cohen. 1998. *Introduction to Logic*, 10th ed. Upper Saddle River, NJ: Prentice-Hall, Inc.
- Fox, John and Subrata Das. 2000. *Safe and Sound: Artificial Intelligence in Hazardous Applications*. Menlo Park, CA: MIT Press.
- Hastings, Arthur C. 1963. *A Reformulation of the Modes of Reasoning in Argumentation*. Evanston, Illinois: Ph.D. Dissertation.
- Hurley, Patrick J. 2000. *A Concise Introduction to Logic*, 7th ed. Belmont, CA: Wadsworth.
- Josephson, John R. and Susan G. Josephson. 1994. *Abductive Inference: Computation, Philosophy, Technology*. New York: Cambridge University Press.
- Kienpointner, Manfred. 1992. *Alltagslogik: Struktur und Funktion von Argumentationsmustern*. Stuttgart: Fromman-Holzboog.
- Layman, C. Stephen. 1999. *The Power of Logic*. Mountain View, CA: Mayfield.
- McGee, Vann. 1985. "A Counterexample to Modus Ponens," *The Journal of Philosophy*, 82, 462-471.
- Peirce, Charles S. 1965. *Collected Papers of Charles Sanders Peirce*, vol. 2. "Elements of Logic," ed. Charles Hartshorne and Paul Weiss. Cambridge, MA: Harvard University Press.
- Perelman, Chaim and Lucie Olbrechts-Tyteca. 1969. *The New Rhetoric*. Notre Dame: University of Notre Dame Press.
- Prakken, Henry. 1997. *Logical Tools for Modelling Legal Arguments*. Dordrecht: Kluwer.
- Reed, Chris and Doug Walton. 2001. "Applications of Argumentation Schemes," *Proceedings of the Fourth OSSA Conference*.
- Reiter, Raymond. 1985. "On Reasoning by Default," *Readings in Knowledge Representation*, ed. Ronald J. Brachman and Hector J. Levesque, pp. 402-410. Los Altos, CA: Morgan Kaufman Publishers.
- Rescher, Nicholas. 1976. *Plausible Reasoning*, Assen: Van Gorcum.
- Sanford, David. 1989. *If P, then Q: Conditionals and the Foundations of Reasoning*, London: Routledge.
- Sidgwick, Alfred. 1893. *The Process of Argument*. London: Adam and Charles Black.
- Toulmin, Stephen. 1958. *The Uses of Argument*. Cambridge: Cambridge University Press.
- Verheij, Bart. 1999. "Automated Argument Assistance for Lawyers" In *The Seventh International Conference on Artificial Intelligence and Law: Proceedings of the Conference*, pp. 43-52. New York: ACM. (Available at <http://www.metajur.unimaas.nl/~bart/>, follow the links.)

- Verheij, Bart. 2000. "Logic, Context and Valid Inference Or: Can There be a Logic of Law?" Available on the same web page as (Verheij, 1999).
- Walton, Douglas. 1996. *Argumentation Schemes for Presumptive Reasoning*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Walton, Douglas. 1996a. *Argument Structure : A Pragmatic Theory*. Toronto: University of Toronto Press.
- Walton, Douglas N. and Erik C. W. Krabbe. 1995. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. Albany: State University of New York Press.
- Wigmore, John H. 1940. *A Treatise on the Anglo-American System of Evidence*, Vol. I (of 10 volumes), 3rd ed. Boston: Little, Brown and Company.

Douglas N. Walton
Department of Philosophy, University of Winnipeg
515 Portage Avenue
Winnipeg, Manitoba, Canada R3B 2E9
Email: d.walton@uwinnipeg.ca
www.uwinnipeg.ca/~walton