What is an Argument?*

1 The Issue

> "I see what your premises are," says the philosopher, "and I see your conclusion. But I just don't see how you get there. I don't see *the argument*."

We hear such comments often. They indicate that there is a notion of "argument" in philosophy in which an argument does not consist just of premises and conclusion; it has additional structure. This distinguishes the notion of argument in philosophy from the technical notion most commonly found in logic texts, where an argument is an ordered pair consisting of the premises and the conclusion. The philosopher's argument is something with more structure, more akin to the logician's notion of *derivation*: a series of statements with intermediate steps providing the transition from premises to conclusion. However, there are substantial differences between arguments and derivations. One difference is in the matter of ordering; philosophical arguments, more often than not, have the conclusion at the beginning rather than at the end, reflecting the fact that individual steps in the argument are validated by later steps, whereas in typical derivations, steps are validated only by previous steps.[1] This difference sounds unimportant until we reflect on the notion of *circular* argument, an important preoccupation of philosophy. A textbook derivation cannot be circular; the ordering prevents it. Circularity begins when you put forth a statement to be validated by something that comes later, and then use that very stating to validate some of the crucial reasoning that comes later. If you try to duplicate this with a formal derivation, it can't be done; you simply fail to produce a derivation. Formal derivations are a means to avoid circularity, not to embody and analyze it. For this reason, among others, they are not what philosophers analyze when philosophers analyze arguments.

Philosophical arguments can also commit other important sins; for example, they can beg the question. Attempts to explain what it is to beg the question that are articulated solely within the terminology of deductive logic have been notoriously unsuccessful. At best they produce no account at all; at worst they give credence to the false claim that all valid arguments are question begging. The result is widespread skepticism about whether the notion of begging the question even makes sense. Yet the notion of question begging is one of the tools of the trade in philosophy in the critical assessment of arguments.

Question begging and circularity are important not just in evaluating arguments; they also apply to explanations, and they even apply to definitions. Definitions can be circular, and they can beg the question. The same with explanations. These are important kinds of evaluation, and we ought to have a good theoretical account of them. We don't.

My goal in this paper is to describe the notion of *argument* as it is used in contemporary philosophy and to describe the methods of evaluation that philosophers use to assess arguments. I won't discuss definitions and explanations here. Part of my task C the easy part C is to extend techniques available within formal logic to evaluate arguments with structures similar to those of formal proofs. The other part of the task is philosophical; we need to get straight on what the issues are. Without some preliminary philosophical ground-clearing, we will be doomed to duplicate well-known inadequacies of informal logic.

This investigation covers some of the same material addressed by traditional texts on informal logic, especially their analyses of the nature of certain fallacies. However, my theoretical approach is different, and so are many of the results. I think that the field of informal logic has been hampered by a lack of theory or perhaps by possession of wrong theory. This can be made right by the development of a better theory.

2 Interpreting Texts *versus* Assessing Arguments

One oddity of argument evaluation is that our evaluations sometimes seem clear and objective, and other times quite unclear and subjective. One reason for this is that evaluating an argumentative text involves both scholarly interpretation and logical assessment.

Arguments originate in texts, written or spoken. In assessing an argumentative text there are two steps: you *interpret* the text, and you *assess* the argument that you have attributed to the text as a result of interpretation. The first step, interpreting the text, is a sophisticated scholarly task. It is typically underdetermined by all available evidence, evidence must often be balanced against counterevidence, and there may be an ineliminable element of subjectivity to it. The second step is the logical task of assessing an argument. This is mostly clear and objective.

We begin with a text, written or spoken, and it is a matter of interpretation whether it contains an argument at all, and, if so, what and where. At the first level of interpretation we locate a bare bones argument, that is, we identify the premises, conclusion, and intermediate steps that are overtly present in the text. I call this the "source argument" or "ur-argument." The ur-argument may be lacking parts that the author expects the reader to fill in, it may be equivocal, and it may be unclear in many other respects. So we usually *interpret* it; we go beyond the ur-argument, seeing it as the overt manifestation of a refined argument, one with a more developed and more definite structure, (or sometimes seeing it as an equally good manifestation of several such developed arguments). This involves filling in steps that are not articulated in the ur-argument, and it involves clarifying meanings, e.g. so as to resolve ambiguity.

When we say that an argument is an *enthymeme* we are speaking of a refined argument *in relation to* a text. We have an enthymeme if the refined argument that we attribute to the text contains one or more steps that are not overtly present in the text. It does not make sense to call either an ur-argument or a refined argument in isolation an enthymeme; what is at issue is a comparison between them.[2] Likewise, when we say that an argument is *equivocal*, we mean that the single ur-argument in the text gives rise naturally to more than one refined argument, related in certain ways.[3] So when we talk about enthymemes and equivocation we are talking about something complex that includes both the text and one or more refined arguments seen to lie in it. Many of our other assessments of arguments, however, are directed entirely towards the refined argument in abstraction from any of its possible sources. I turn next to a description of what a refined argument is, and what it means for a refined argument to be successful.

3 Refined Arguments

Suppose that we have fully interpreted an ur-argument from a written or spoken text. The interpretative process then must have yielded three things: a setting, a target proposition (the intended conclusion), and a reasoning structure propounded within that setting as a means of reaching the target:

REFINED ARGUMENT		
SETTING The Rules and Assumptions	TARGET The Proposition to be Established	REASONING STRUCTURE The Steps in the Argument and Relations among them

3.1 Settings

A setting includes at least the following:

A set of statements that are taken for granted: the set of assumptions.

A set of inference rules that are taken as acceptable for purposes of reasoning.

These should be familiar to us. When G. E. Moore (1939)[4] argues that there are external objects, he argues within a setting in which he takes for granted that there is a hand in front of his face. When Descartes argues that there are external objects, he does not take for granted that there is a hand in front of his face. Moore's assumption is part of the setting within which his arguing takes place. Descartes cannot argue as Moore does, because his assumptions are different. (Descartes and Moore are unusually explicit in telling us what their assumptions are.)

Settings also include assumptions about what rules are acceptable. Suppose I argue that universals are linguistic, since they have to be either linguistic or in the world, and they are nowhere in the world. If I argue like this, I take for granted the law of disjunctive syllogism: from A_B and $\_A$ infer B. But when certain logicians argue about non-classical logic, they do not assume that disjunctive syllogism is a valid principle. In some cases, this inference isn't available to them as an assumed principle; it isn't part of their setting.

So settings contain both assumed propositions and assumed principles of inference.

3.2 Targets

It is part of our notion of argument that it has a *goal*, which is to establish some particular proposition. That is all that I mean by a "target" C it is a proposition to be validated.

Redundancy:

An argument is *redundant* if its target is already assumed in the setting, and otherwise *nonredundant*.

A redundant argument pursues such a pointless goal that serious texts are rarely interpreted as containing them. But the possibility of their existence is theoretically important, if only to distinguish redundancy from other defects. If the redundancy is apparent, the most efficient reasoning structure would consist of announcing the conclusion and identifying it as something we assume. In fact, we often do this, and we even call it arguing; we say "so-and-so argued that *P*" when *P* was one of the assumptions that so-and-so articulated and insisted upon, but gave no further argument for. So we might as well include *insisting on something that is being assumed* as a (redundant) form of argument.

3.3 Reasoning Structures

A reasoning structure is similar to a derivation in logic; it consists of a sequence of statements[5] that is meant to reach a target within a setting, with the statements annotated to identify them as premises, conclusion, or as inferred from others. Most of the detailed work of the theory of applied logic is to describe what a reasoning structure is, and to see how it bears on the success of the argument.

For simplicity here, I ignore arguments that contain subproofs. With that simplification we can construe a reasoning structure as a sequence of statements, with each member of the sequence (each "step") either identified as a "premise" or identified as directly inferred from some set of other steps in the sequence. In addition, one of the steps is identified as the conclusion.

> A *reasoning structure* is a sequence of "steps," in which one step is categorized as the conclusion, and in which each step is categorized as a premise or is as being inferred from an identified set of other steps. (Nothing is categorized both as a premise and an inferred step.)

Notice that this characterization imposes no constraint on the order of occurrence of the steps; in particular, steps may be inferred from both earlier and later steps. The lack of constraint on order permits reasoning structures to be circular:

> A *circular* reasoning structure is one in which the inference dependencies can be traced from some step back to itself.[6]

I have referred to steps in an argument being "categorized" as premises or as inferences from other steps. This is a product of interpretation. That is, I assume that identifying a fully interpreted argument involves identifying each of its steps, and identifying the status of each step; this involves identifying it as a premise of the argument in question, or, if it is not a premise, identifying which of the other steps supposedly validate it. If we do not know this much, then we do not know *what* the interpreted argument is.[7]

3.4 The Success of Arguments, Conceived as Tasks

An argument is a reasoning structure *in* a setting *with* a target. We tend to speak as if the reasoning structure itself is the argument, but it is essential to the various ways in which we evaluate an argument that we presuppose a setting and a target. I claim also that when we evaluate an argument we conceive of it as the carrying out of a *task*: to arrive at a specified target within a given setting by certain means. Such an attempt may or may not be successful in its own terms. The conditions of such success are these:

A *successful* argument is one in which

> $ Every premise is among the statements assumed in the setting.

$ Every inference is in accordance with a principle of inference assumed in the setting.
$ The conclusion is the target identified in the goal.
$ The reasoning structure is noncircular.
$ There is no infinite regress of justifications for any step.[8]

This is meant to be a *description* of what contemporary philosophers take to be successful arguments, in the sense that this is what philosophers presuppose when they evaluate arguments.

4 Examples of Arguments

In discussing the theory, it will help to fix on some specific rules of inference to take for granted in giving illustrations. I assume that the principles of classical logic give us one example (one among many) of what might be contained in the rules in a setting. Suppose that our statements are couched entirely in the terminology of the first order predicate calculus, and let our set of rules of inference R be those that admit as immediate inferences the traditional classical ones (such as *modus ponens*, *disjunctive syllogism*, *universal instantiation*) listed in logic texts, with a list sufficiently broad so that we have a complete system without subproofs. Then, as one would hope:[9]

S is provable from a set of sentences Γ in the predicate calculus iff
there is an *argument (as defined above)*
which has $\Gamma \Updownarrow R$ as its *setting (as defined above)*
and has S as its *target (as defined above)*
and is *successful (as defined above)*.

Until further notice, I will assume that each setting assumes some set of classical rules of this sort.[10]

If a reasoning structure is noncircular, you can rearrange its steps so that inferred steps are inferred entirely from previous steps. So any successful argument with a finite number of steps can be rearranged into a proof of the standard sort encountered in logic texts. (This can never be done for circular arguments.) The ordering of proofs, then, can be seen as a device for avoiding circularity.

Here are some examples of reasoning structures that further illustrate the function of a setting. `⊲' marks premises and `#' marks the conclusion, and lines and arrows indicate the inferences:

EXAMPLE 1: The reasoning structure of a successful argument in a setting similar to that of G. E. Moore's discussion of external objects:[11]

```
◁    There is a hand in front of my face.  --|
                                             |
◁    Hands are material objects.           --|
                                             |
#    There are material objects.      <------|
```

I am imagining here a setting in which a momentary doubt has arisen concerning the abstract question whether there are material objects, but in a situation in which nothing has happened to raise Cartesian doubts. Your ordinary beliefs have not been called into question in any way; you are just momentarily disconcerted by the abstract issue: are there material things? Marshalling your mental resources, you articulate the above reasoning, and set your mind at rest. You have given a successful argument.

EXAMPLE 2: The reasoning structure of an unsuccessful argument in Descartes' setting (the first premise is not among the assumptions):

```
◁    There is a hand in front of my face.  --|
                                             |
◁    Hands are material objects.           --|
                                             |
#    There are material objects.      <------|
```

In Descartes' setting the same reasoning process is not successful, because the assumptions are different. A similar point applies to Berkeley's setting:

EXAMPLE 3: The reasoning structure of an unsuccessful argument in Berkeley's setting (the second premise is not among the assumptions):

```
◁    There is a hand in front of my face.  --|
                                             |
◁    Hands are material objects.           --|
                                             |
#    There are material objects.      <------|
```

5 Fallacies

I am giving something like a transcendental argument for arguments having a certain kind of structure: this is the structure arguments need to have in order for us to assess them in the ways in which we do. It is thus appropriate to briefly survey the traditional topic of fallacies[12] as a compendium of ways in which we evaluate arguments.

It is common in informal logic to divide fallacies into

Fallacies of ambiguity,
Formal fallacies, and
Fallacies of relevance.

These apply to different components of arguments:

Fallacies of ambiguity	C	Ur-arguments + reasoning structures.
Formal fallacies	C	Reasoning Structures.
Fallacies of relevance	C	Refined arguments.

5.1 Fallacies of Ambiguity

These are the traditional fallacies of ambiguity: *equivocation*, *amphiboly*, *accent*, *composition and division*, and so on.[13] They pertain to the relation between ur-argument and refined argument. They are tricky to define, but we usually have little difficulty in recognizing them, and I will not discuss them here.

5.2 Formal Fallacies

Formal fallacies include things like affirming the consequent. Three of the conditions for success given above bear on formal fallacies. An argument that fails to meet the second condition of success (that each inference accords with a rule in the setting) commits a *non sequitur* if the step appeals to a rule in the setting but the step does not actually follow by that rule.[14] Failures of the conditions of *noncircularity* and *lack of infinite regress* are also formal fallacies.

5.3 Fallacies of Relevance

Certain of the traditional "fallacies of relevance" apply at the level of refined argument and nowhere else. They correspond in part to the conditions for successful arguments given above.

Begging the Question: The first condition for success for an argument given above is that each premise be among the things assumed in the setting. If a sentence used as a premise is not among the assumptions in the setting, Whately[15] calls this the fallacy of "Premise unduly assumed." If that premise is identical with the conclusion itself, or equivalent to it, this is the traditional fallacy of *begging the question*. At least, this is the definition of begging the question that is given in virtually all logic texts that address this issue. But this is not how contemporary philosophers use the term `beg the question'. In modern usage, almost any "undue assumption" is called begging the question if the assumption used is crucial to the argument.[16] For example, if Descartes had been trying to prove that material things exist, and if he used as a premise that his desk exists, then he would be accused of begging the question, even though this is not assuming his sought-for conclusion, "that there are material things". Modern usage thus departs from textbook definitions that respect the etymology of the term.[17] For modern purposes, violation of the first condition for success of an argument (when the step in question is crucial to the argument) can be taken to be a pretty good characterization of the logical fallacy of begging the question:

> *Begging the Question:* An argument begs the question (in the logical sense) if one of its premises is not among the assumptions of the setting.

(I argue below that there is a distinct, non-logical notion of "begging the question" that is used in the assessment of settings themselves.)

The principal difficulty in deciding whether an argument begs the question is deciding what the full argument is, principally by deciding what the setting is. Given a text, it may be difficult to decide whether the author is proposing to argue from certain assumptions that include the step in question -- in which case no question is begged -- or whether the author is not including that step among the things being taken for granted. This is a question of interpretation, not of logical assessment, and in particular cases it may be difficult or even impossible to settle. But there is no logical puzzle here, nor any gap in logical theory. Once the setting is clear, the logical assessment is straightforward.

Another lack of success is this:

> *Ignoratio Elenchi:* An argument that fails to meet the third condition of success argues to a conclusion different from the point at issue.

This seems to be the most trivial and easily detectible of errors, but I agree with Schopenhauer[18] that this is the fallacy most commonly committed, and also the most effective means of actually deceiving people (including the person giving the argument).

> *Petitio Principii:* This is the specific version of begging the question in which the conclusion itself is used as a premise. A nonredundant argument that uses the conclusion as a premise begs the question if the conclusion is the target; otherwise it is an *ignoratio elenchi*.

It is apparent that distinguishing these fallacies relevant to full arguments from one another, or even ascertaining whether any of them actually occur, depends essentially on accurately identifying the setting, on identifying what is and what is not being assumed. This is one reason why informal logic textbooks sometimes seem to lack coherence; the exercises they set for students assume that fallaciousness is an inherent property of a reasoning structure apart from the setting. So they use short passages from texts, with the assignment being to decide whether or not the passage begs the question. But the reader does not have sufficient information to identify the setting. This makes it impossible to tell for sure what fallacy is being committed, if any, and mistakenly gives credence to the view that there is no subject matter here susceptible of serious investigation.[19]

6 Circular Reasoning

6.1 Circular Arguments

In popular speech almost any fallacy is labelled "circular" reasoning. I focus in this section on those special cases in which the notion of a "circle" in reasoning is a helpful metaphor in assessing logical structure.

The traditional literature in logic discusses two quite different notions of logical circularity. The first is pertinent to the notion of circularity introduced above. In this sense, an argument is circular if the inferences in its reasoning structure lead from some step back to that very same step. An example is the following:

> Mary will win the election (CONCLUSION). Because, if she runs, she'll win (PREMISE), and she'll run. And she will run because she'll run iff Bill doesn't, and (PREMISE) he won't run. The reason she'll run iff he won't is that (PREMISE) if either of them run, the other won't. And at least one of them will run, because she will run (ESTABLISHED ABOVE).

Sample formalization:

```
#    Mary will win.                    <--|
                                          |
◁    If Mary runs, she'll win.          --|
                                          |
     Mary will run.                     --|<-|       -|
                                             |        |
     Mary will run iff Bill doesn't.       -| <-|     |
                                             |  |     |
◁    Bill won't run.                       -|   |     |
                                                |     |
◁    If Mary or Bill runs, the other won't. --|     |
                                                |     |
     Mary or Bill will run.                   --| <-|
```

In this argument, the third, fourth, and last steps have justification paths that jointly form a "circle," so each has a justification path that leads back to itself. There is no *non sequitur* in this argument, and there is no reason to see any premise as begging the question, so the *only* fallacy inherent in an argument of this form is due to circular reasoning.

Here is a difference in detail from most of the informal logic literature. There, `circularity' is usually defined as an argument's using its own conclusion (or something logically equivalent to its own conclusion) as one of its premises.[20] But this cannot be the full story, for on this account the displayed argument is not fallacious at all! Yet it is clearly a bad argument, and a paradigm of what we normally call circularity. (The argument *is* circular on the account of circularity given above.)

It is important to keep in mind the difference between inferring a statement from itself and inferring a step from itself. The latter is the essence of circular reasoning. But the former is not; it may or may not be perfectly legitimate. For example, any repetition of a statement previously invoked in an argument is "inferring" a statement from itself, but this is hardly circular. Thinking so easily leads to popular mistaken critiques of deductive logic. If an inference from *A* to *A* is a fallacy, can inferring *A & A* from *A* be much better? And then how about inferring *A & B* from *A* and *B*? And then . . . . This slippery slope strongly suggests that the essence of circularity is validity, and then we have John Stuart Mill breathing down our necks with his infamous view that all deductive reasoning begs the question.[21] Inferring a statement from itself is not fallacious; reaching a *step* in your reasoning that is directly or indirectly based on that very step is fallacious.

Perhaps inferring *A* from itself is not in general a fallacy, but sometimes it seems to be. For example, what about *concluding A* from *A*? That is, what about inferring *A as a conclusion* from *A* used *as a premise*. Surely there's something fallacious about that, isn't there? What about the famous case that disturbs students who first undertake to study logic, the purported validity (perhaps even soundness) of:[22]

```
      A       ---|
  ---------      |
     — A    <--|
```

On the account of argument given here, an argument of this form is *not* circular. But it inevitably fails (if it is not redundant), and we know this from its form. We can't specify why it fails until we know more about the context, but it will fail. In particular, any argument of this form will fall into one or more of the following categories:

Redundant	C if its conclusion is among the assumptions.
Question-begging	C if its premise is not among the assumptions.
An *Ignoratio Elenchi*	C if its conclusion is not the target.

Any of these options are possible, and arguments of the given form inevitably satisfy one or more of them. So no argument of this form can be both successful and nonredundant. We might want to call this "circularity," but it is important to distinguish this rather special kind of circularity from the sort defined above. This illustrates the importance of keeping in mind the full argument, not just its reasoning structure, or its form, when considering even such formal notions as circularity. This importance is illustrated further by the second kind of circularity.

6.2 Circularity of Combinations of Arguments

Many discussions of circularity in the informal logic tradition address a different phenomenon: circularity due to the interaction of *two or more* arguments. The popular example is

> God exists, because it says so in the Bible.
> We can trust what the Bible says, because it is the word of God.

This example is messy to deal with because its parts are enthymemes; neither argument is successful as stated even granting the stated premises. But the point is clear enough: some cases of circularity consist of two arguments, where the conclusion of each serves as a premise of the other. A simple case is:

A	B
A □ B	A □ B
———	———
— B	— A

Neither of these forms of argument is automatically objectionable in itself, yet there is something disconcerting about their combination. The "circularity" of the combination comes from the fact that if we link the conclusion of each with the first premise of the other, we get an inference path that traces certain steps back to themselves:

```
|--   A    <- --|  |-->   B   --|
|               |  /            |
|--   A □ B   |-|-|   A □ B --|
|             |  |              |
|--> — B   --|  |---- — A <--/
```

But a "circular" diagram, by itself, tells us nothing; why should this pattern of multi-argument circularity be fallacious?

In fact, there are two quite different sorts of examples to consider. The first is the case in which both arguments are given in the same setting. In this case, the circular path shown above is heuristic, because it illustrates the following easily established fact about multi-argument circularity:

> Whenever there are two arguments *in the same setting*, each using the conclusion of the other as one of its own premises, it is impossible that both arguments be both successful and nonredundant.

Multi-argument circularity of this sort in a common setting guarantees that something will be wrong with one or both of them, though we cannot tell in detail what will be wrong without knowing more about the setting.

What, then, about two arguments so related but in different settings? This isn't necessarily a fallacy at all. An example is the politician who tells the farmers on Tuesday that her party's platform is good because it advocates farm subsidies, and tells a group of committed party supporters on Wednesday that they should advocate farm subsidies because the platform calls for it. On Tuesday she assumes that farm subsidies are good and takes support for the platform as her target; on Wednesday she assumes that support for the platform is good and takes farm subsidies as her target. This may or may not be duplicitous, depending on circumstances, and it may or may not mislead, but it is hard to see it as an automatic fault of reasoning. The error comes only if you argue from different assumptions on different occasions, and then pretend that you haven't.[23]

6.3 Arguing in Big Circles

It has become popular over the last decade or two for philosophers to say that it is all right to argue in a circle if the circle is big enough. (A big enough circle is a "virtuous" circle.) This is a puzzling claim, and deserves scrutiny. It generally occurs in a discussion of epistemology, and seems to be applied to global sets of general beliefs about the world. The point seems to be that if we can argue in a circle that involves many of these, this somehow verifies all of them. But how can this be? Suppose that the indicated inferences in an argument of the following form are all valid:

```
A1   --|             <-|
       |               |
A2   <-|  --|          |
            |          |
A3        <-|          |
...                    |
...              --|   |
                   |   |
An               <-|  --|
```

Then, in classical logic, so are these:

```
_An            --| <-|
                 |   |
...            <-|   |
...                  |
                     |
_A3       --|        |
            |        |
_A2  --|  <-|        |
       |             |
_A1  <-|           --|
```

But if a "big circle" argument of the former kind validates all of its constituents, so should one of the latter kind. There must be more to the story than this.

I suggest that part of it is this. The circularity is meant to be a multi-argument one, where each argument has a setting that is at least partly epistemically independent of each of the others. Each setting is based on different assumptions, but on ones that we do in fact accept, and each argument is (by hypothesis) successful. The point of the completed "circle" is that it shows that any one of the settings that we accept leads to any of the others, so we cannot abandon any of them without abandoning all of them. This then renders each more secure. The reverse "big circle" *would* show the same thing for each of its constituents, if we accepted them, but since we accept none of them the point is of no particular relevance.[24]

7 Settings

If the theory given so far is correct, then we have an account of what arguments need to be like in order to allow for the *logical* assessments of them that we make. But we also assess arguments as regards their epistemological status. Epistemological assessments are assessments of their settings, independent of their reasoning structures.

As I have construed arguments, a person can argue for anything at all on the basis of any assumptions at all, and there are few limitations on which such arguments can be successful. For example, one can argue successfully from totally implausible premises. Isn't something missing here? Further, whether an argument is successful or not is a relatively straightforward and objective matter. But logic is often said to be a normative, prescriptive enterprise. Where does the normativity get in? These two questions are linked: isn't there some way to say that you *shouldn't* argue on the basis of implausible premises?

I see logic itself as objective and non-normative. It is the *uses* of argument that give rise to normativity, and the normative issue has primarily to do with the choice of settings.

The choice of settings gets discussed rarely in logic texts, more frequently in introduction to philosophy texts and in philosophical journals, usually under the (ubiquitous) heading "begging the question".[25] This is a different notion of begging the question than the logical one discussed above. This version is:

> You beg the question if you use as a premise a proposition that is as dubious as the conclusion, or you use as a premise a proposition that is not knowable independent of the conclusion, or . . .

It is clear from this formulation that the status of such a principle is to *constrain appropriate settings*. The proposals come from a framework in which argument is seen as embodying reasoning within a process of inquiry or of belief formation. The enterprise is then an epistemological one, not just a logical one, and the constraint comes from the epistemological application of the argument; the argument cannot do its epistemological work (no matter how successful it is logically) because its choice of setting begs the question.

A setting is regarded as epistemologically inappropriate unless the propositions assumed in the setting are better known (or more readily knowable, or less subject to doubt, . . .) than the target. Further, the principles of inference assumed should be ones that potentially preserve knowledge; that is, they should be principles such that if you know certain propositions and an inference rule leads from them to another proposition then reflection on this fact should enable you to know that other proposition too, and these rules should be less dubitable than the point at issue.[26]

These considerations *are* normative. They use the notions of knowability or dubitability, which are normative. For example, it does no good to define an appropriate setting as one in which someone might believe the assumptions in the setting while doubting the target, because this is too capricious; "irrational" people may doubt or believe any combination of propositions. You must appeal instead to the beliefs and doubts of a *reasonable* person who has carefully considered the issues; you require that a *reasonable person* might believe the assumptions while doubting the target. Thus normativity creeps in.

It is important to separate out these epistemological constraints on settings and targets from the logical aspects of argument, for at least three reasons. First, people patently do give arguments that are regarded as successful even when no modern philosopher thinks that they achieve any such epistemological purpose. (Plato gives lots of these; so does Aristotle, and so on.) Second, it is important in our assessment of arguments to distinguish violating these principles from violating the more purely logical ones I have catalogued. As a simple illustration, an argument may beg a question in the epistemological sense by arguing from premises some of which are dubitable without begging the question in a logical sense at all. The argument is impeccably implemented, it just happens to be an epistemologically pointless one. Unless we can disentangle these points we will not be able to teach either logic or epistemology with understanding. A third reason for distinguishing the epistemological considerations is that on close examination it may turn out that there are no such considerations. It is notorious that constraints such as the one that requires the premises to be less dubitable than the conclusion sound fine in the abstract but are almost impossible to recognize in the concrete. Dubitability is an equal opportunity assessor, which tends to treat propositions alike; and when it does distinguish them, valid consequences of principles tend to have *less* dubitability than the principles themselves, frustrating the idea that we need to validate the more dubitable in terms of the less.

Still, something like the standard view must be correct. At least the following is true: An argument serves an epistemological purpose when, *on the spot*, it gets you from something (or some things) *you* didn't doubt to something you did. This may or may not be because what you doubted at the time had a lower *degree of dubitability* than what you didn't doubt. In any event, this sort of consideration is clearly a constraint on settings, not on reasoning structures *per se*.

In summary, I see arguments as hemmed in between scholarly interpretation on the one end and epistemological application on the other:

Scholarly Interpretation	$\longrightarrow$	The Argument	$\longleftarrow$	Epistemological Purpose

The scholarly interpretation is subject to subtle and complex constraints, some of them normative. And so is the epistemological purpose. But the argument itself, hemmed in between these, can be both delimited and assessed in a clear and non-normative fashion. This is the locus of logic proper.[27]

Terence Parsons

Notes

* I am indebted to many people for discussions of these views, but especially to Peter Woodruff and Penelope Maddy.

---

1. Some derivations, such as those in Kalish & Montague, are constructed so that formulas are validated by subsequent subderivations (Kalish, Donald and Richard Montague, *Logic: Techniques of Formal Reasoning*, Harcourt, Brace, and World, New York, 1964). The point I am making applies to these too, though it needs a slightly more complicated formulation.

2. Some have wanted to apply the notion of enthymeme to ur-arguments alone, by asking whether the argument *needs* additional steps filled in in order to be successful. Such tests run the danger of turning all invalid ur-arguments into enthymematic valid ones. This goes against the plausible assumption that people sometimes argue incorrectly. If there is a test for determining whether there exists a missing premise, the test is a factor in scholarly interpretation, not in logical assessment.

3. By calling an argument equivocal we do *not* mean merely that the ur-argument contains an ambiguity. Most arguments in English contain ambiguous words that do not bother us at all, and that do not lead to the charge of equivocation.

4. I refer here to Moore, G. E., "Proof of an External World," *Proceedings of the British Academy* 25, 1939, 273-300. This is different from Moore's argument in "A Defense of Common Sense," in J.H. Muirhead (ed.) *Contemporary British Philosophy* (Second Series), 1925, Macmillan, London, 193-223.

5. It is not clear whether there is anything in the nature of argumentation that requires the steps in an argument even to have an order. But stated arguments always consist of statements produced in some particular order, and for this reason I view reasoning structures as sequences of statements together with certain relations among them. The irrelevance of order should be a product of investigation, not an initial datum.

6. Define a "dependency path" for a step *s* as any sequence whose first member is *s*, and such that any member of the sequence that is not a premise is followed immediately in the sequence by one of the steps from which that member is immediately inferred in the reasoning structure (and any premise is followed by nothing; i.e. it terminates the sequence). Circularity results if some step has a dependency path that includes the step itself as a member other than the first.

7. To take account of conditional and indirect reasoning we need to expand this framework. First, expand the notion of a reasoning structure to include subarguments:

> A *reasoning structure* is a finite sequence of steps which are either statements or subarguments, in which one step that contains a statement is marked as the conclusion and in which each step in the sequence that is a statement is either marked as a premise or is marked as being inferred from an identified set of other steps in the sequence.

For this new definition, we need an account of subarguments:

> A *subargument* of a reasoning structure is a finite sequence of statements in which the first step is marked as the hypothesis of the subargument and every other step is marked as a premise or as being inferred from an identified set of other steps in the sequence or in the reasoning structure itself.

This permits subarguments of reasoning structures. In order to get subarguments of subarguments, extend the definitions in the obvious way.

I suppose that the hypothesis of a subargument is always stated first in the subargument, and that all of the steps of the subargument are in one contiguous series, uninterrupted by steps of the main argument. These two conditions seem to be met already in texts, the only apparent exceptions being enthymemes.

The notion of circularity given below in the text also needs refining, by extending the notion of inference dependency to include inferences involving subarguments. We do this by assuming that any statement that is directly inferred from a subargument is dependent on those steps outside the subargument that the steps in the subargument are dependent on (assuming that hypotheses of subproofs are not dependent on anything).

Lastly, we need to add to the conditions for success of an argument (given below in the text) that

> *The conclusion does not depend on any hypothesis*,

where again we analyze dependency by tracing lines of inference dependency. (This condition does not correspond to any of the traditional fallacies; this is no surprise since they were developed in times when arguments with subproofs were not recognized.)

Here is an example of a conditional proof to illustrate these notions, using `&' to indicate hypotheses and using indentation to indicate subproofs:

```
◁     A _ B      --|
                     |
◁     B __ C        |  --|
                     |    |
&          _A    --|    |    |
                     |    |    |
           B     <-|  --|    |---|
                          |    |   |
           C        <--|    |   |
                                   |
#     _A __ C              <------|
```
This is an argument with four steps, one of which is a subproof. The last line is inferred from the third step, that is, *from the subproof*. It depends on all the steps that the steps of the subproof depend on from outside the subproof, namely, the first two steps, which are premises. Tracing back dependencies yields no circularities and no dependency on hypotheses. (If the next to last line (`C') had been marked as the conclusion, it would violate the condition for success that prohibits dependency on hypotheses, since that line depends on the hypothesis of the subproof.)

8. As above, define a dependency path for a step *s* as any sequence whose first member is *s*, and such that any member that is not a premise is followed by one of the steps from which that member is immediately inferred in the reasoning structure (and any premise terminates the sequence). There is a regress of justifications for *s* if there is an infinitely long dependency path for *s* that contains no repetitions of steps. (If we drop the condition that there be no repetitions, then circularity becomes a special case of regress. Sometimes we distinguish circularity from regress and sometimes we see circularity as a special case of regress. I have given the definitions in a form that allows us to distinguish them.)

It is clear that infinite regress must be ruled out. Here is a text that is naturally interpreted as containing an argument proving that 1 is irrational; the argument has no flaws except regress:

> *"If a number is rational, so is that number plus 1. So 1 is irrational, because 2 is irrational. And 2 is irrational because 3 is. And so on."*

It would be rash to avoid the issue of regress altogether by requiring that arguments be finite. Even if a text containing an argument must, by the nature of a text, be finite, this would not extend to refined arguments embodied in the text. Imagine a setting that contains among its assumed rules of inference the *omega rule*:

> From all of `*Pn*' where *n* is a positive integer, infer `Every positive integer is *P*'.

Suppose also that the setting contains the assumptions that 1 is finite, and that if a number is finite, so is its successor. Then the following text can be interpreted as containing a successful argument with an infinite number of steps:

> *"1 is finite. So 2 is finite. And so on... So every positive integer is finite."*

9. If S is provable from Γ then there is a classical derivation of S from Γ, and this derivation is already the reasoning structure of a successful argument with $\Gamma \Updownarrow R$ as the setting and S the target. For the opposite direction, if there is a successful argument of the indicated sort with a finite reasoning structure, then its reasoning structure is non-circular, and it can be rearranged into a finite sequence of steps with each step inferred from previous steps. With proper labelling, this will be a classical derivation. If the reasoning structure is not finite, it can still be rearranged so that all inferences are from earlier steps, except that it may be infinitely long back to front (that is, starting with the last step you may be able to proceed backwards an infinite number of steps). In this case, consider the argument you get by including the conclusion, and any steps from which it is inferred, and any steps from which those steps are inferred, and so on. By the nature of the rules in R we may discard any other steps and still have a successful argument. The resulting pruned argument cannot be infinite, because otherwise it would violate the constraint against a regress of justifications for the conclusion. So the pruned argument, properly labelled, is a classical derivation.

10. The familiar rules of deductive logic are all context free, in the sense that whether they apply depends only on the presence of sentences cited in the reasoning structure. But some much-discussed rules are not like this. Consider the "statistical syllogism": *If most A's are B and if x is an A then x is a B*. The problem with this rule is that if you know its premises, but you also know that all C's are not B, and you know that x is a C, you would be a fool to apply the statistical inference rule. So the rule is usually presented with a qualification:

> If you know that most A's are B, and you know that x is an A, then you can infer that x is a B *provided that you have no other relevant information about x*.

This type of rule easily fits into the framework under discussion. Instead of alluding to what you know, you refer to what is present in the setting:

> If most A's are B, and x is an A, then infer that x is a B *provided that no other relevant information about x is present in the setting*.

I do not mean to suggest that this is an easy rule to formulate, since a good formulation would require getting clear on what it means for "information" to be "present" in a set of propositions. The point is rather that the context sensitivity of this sort of principle of inference is allowed for in the characterization of *argument* under discussion.

11. Moore's own argument (from Moore 1939, *op. cit.*; *not* Moore 1925, *op. cit.*) is a "proof" of the existence of *external* objects, whereas the illustration I give is an argument for the existence of *material* objects. So this is not Moore's argument; it is a different one patterned after his.

12. An excellent account of the tradition is found in Hamblin, C. I., *Fallacies*, Vale Press, Newport News, VA., 1970. In my survey I concentrate on the Aristotelian tradition, ignoring entirely the "fallacies" invented by Locke: appeal to pity, appeal to force, etc. These are interesting, but they have little bearing on philosophers' evaluations of arguments. Hamblin's book has inspired a substantial literature that investigates fallacies defined within two-person dialogues. I ignore that literature because of my focus on argument (which I do not see as dialogue).

13. In the tradition on fallacies, \`equivocation' means ambiguity of a single word and \`amphiboly' means a certain kind of ambiguity in grammatical structure. The fallacies of composition and division, in their traditional forms, include fallacies based on the "distributive/collective" ambiguity, such as the ambiguity in \`The boards are heavy' and the ambiguity in \`A person is necessarily two-legged'. The fallacy of accent is variously interpreted, and usually refers nowadays to ambiguities that can be resolved by stress. Most of what we now call "scope ambiguities" would be classified as cases of amphiboly or composition/division.

14. There are perhaps two ways for an inference step to fail. It can fail because the step is an imperfect implementation of the rule appealed to in the step, when the rule is one of those assumed in the setting. This is a *non sequitur*. But perhaps the step can also fail because it is a perfect implementation of a rule of inference that is not assumed in the setting. In this case, the argument *begs the question*. The fact that we distinguish these two cases in practice suggests that we might want to include the identification of the inference rule as part of the reasoning structure.

15. Page 183 of Whately, Richard, *Elements of Logic* (reprint of the 1827 edition), Scholars' Facsimiles & Reprints, Delmar, N.Y., 1975.

16. As a first approximation, a step is crucial to an argument if the argument would be successful were the step acceptable, and such that there is no way to preserve the success of the argument by erasing that step or a group of steps containing it.

17. If begging the question were limited to cases in which one presumes the conclusion itself, or something equivalent to it, then it would be easy to construct terrible arguments that commit no named fallacy at all. So the definition given here of begging the question reflects an advantage of the modern philosophers' usage over the textbook definitions.

The identification of the fallacy as "begging the question" originates with Aristotle, in *Sophistical Refutations*, 166b.25, in the context of a discussion of "arguments used in competitions and contests" (165b.12). In such a context premises are granted by the opponent, and it is said to be a fallacy to win such a contest by relying on the opponent's foolishly granting the point at issue. (Clearly, this should include granting damaging points *other* than the conclusion itself.) Aristotle also uses the same terminology (`beg the question') in the *Prior Analytics* (64b 33), where he says it is the attempt to prove what is not self-evident by means of itself. This is *not* the same fallacy; I discuss this other "fallacy" in section 7.

18. Schopenhauer, Artur, "The Art of Controversy," in his *Essays from the Parerga and Paralipomena*, trans. by T. Bailey Saunders, London, Allen and Unwin, 1951.

19. A small number of texts partially avoid this problem by giving long passages that provide ample clues to the setting. This still requires the students to merge two tasks: interpret the passage and assess the logic of what you have interpreted. This is a valuable exercise, but it still mixes interpretation with assessment. This is easy to change. You need only assign a text to analyze, and ask the students to (a) interpret the text so that it begs the question, and *then* (b) interpret the text in another way so that it does not beg the question, and assess the argument on that interpretation. This clarifies, rather than obscuring, the dual roles of interpretation and assessment.

20. Yes, this is the same as the traditional definition as begging the question, and some authors even point this out.

21. There is a tradition in nineteenth century logic texts of agreeing with Mill that any valid one-premise argument is circular, but claiming that this is not true of two-premise arguments.

22. In this discussion I assume that the appearance of *A* above the line indicates that it is marked as a premise, that the line indicates an inference, that the "therefore" sign indicates the conclusion, and that the argument has no more structure than this.

23. William Hamilton accuses Plato of effectively concealing a multi-argument circularity by means of intervening material:

> "Plato, in his *Phaedo*, demonstrates the immortality of the soul from its simplicity; and, in the *Republic*, he demonstrates its simplicity from its immortality." (Hamilton, William, *Lectures on Logic* Vol. II. William Blackwood and Sons, Edinburgh and London, 1874, 55. This is also called Volume IV of *Lectures on Metaphysics and Logic*.)

Socrates would enjoy making a case that no fallacy was committed because the arguments actually proceeded from different assumptions; that is, the settings were not the same.

24. This point is made, I think, in Cohen, Morris and Nagel, Ernest, *An Introduction to Logic and Scientific Method*, Routledge and Kegan Paul, London, 1934, 379: "...there is a difference between a circle consisting of a small number of propositions, from which we can escape by denying them all or setting up their contradictories, and the circle of theoretical science and human observation, which is so wide that we cannot set up any alternative to it."

25. This philosophical literature on begging the question is too large for me to do justice to here.

26. Perhaps it is also necessary to add that the rules should not allow the premises to be *undermined* by showing that they entail something dubious, for otherwise the argument might lead one to retract the initial assumptions instead of augmenting them with the conclusion. Without such a requirement it is not clear how argument can be trusted to achieve its epistemological purpose of increasing knowledge. But it is doubtful whether there are any such rules; the traditional rules of formal logic contain no such guarantee. The role of argument in epistemology is mysterious.

27. (This note abbreviates a much longer appendix that has been removed to satisfy journal constraints.) The structures of arguments proposed here probably need enhancing in order to properly accomodate *refutations*, and the technique of *reductio ad absurdum* when classical logic is not presupposed. Here is why. You *refute* a view (an assumption or set of assumptions) by producing an argument from it that concludes with something "absurd" or unacceptable. But what is unacceptable? In classical logic, the negation of any proposition assumed in the setting is unacceptable. If you can derive one of these, you can conjoin the derived proposition with its negation assumed in the setting to get an explicit contradiction, and contradictions are automatically "absurd". This is the traditional technique of *reductio ad absurdum*. But what if your opponent is a maverick who believes that some contradictions are true? Can you still refute such a person by deriving a contradiction from his/her

view? Well, you can from your point of view, of course: you merely show that the maverick's view leads to contradiction. But this leaves some observers dissatisfied; they sense that your opponent has a coherent position that has been left unscathed by your counterargument. What is needed is to derive something from your opponent's view that is unacceptable to your opponent! This may be impossible, of course, if your opponent's view is to accept absolutely everything. This is *omniletheism*; it is a coherent position, but is not particularly interesting because of the lack of distinctions it makes. For this reason it is pragmatically pointless. So suppose instead that your opponent thinks that some contradictions are true (the ones involving the semantic paradoxes, say) and some aren't; the modest *dialetheist*. (This is the view of Graham Priest, *In Contradiction*, Martinus Nijhoff, Dordrecht, 1987.) Surely this is a potentially coherent position, and one that cannot be refuted merely by showing that it leads to contradictions. It can be refuted only by showing that it leads to those contradictions that the modest dialetheist denies.

Let us then refine our notion of setting to include both assumptions and counterassumptions. Then a "refutation" is a successful argument that concludes with something unacceptable in the setting: a *counterassumption*. In classical settings, the counterassumptions are exactly the negations of the assumptions, but in nonclassical settings this need not be so.

Counterassumptions may also be needed to make sense of more conservative points of view. Consider for example those who think that the liar sentence lacks truth value. This is probably the commonest view around, but it is tricky to maintain. Suppose that *s* is a liar sentence; it *is* a sentence that says that it itself is not true:

s = \`*s* is not true'

Suppose we then adopt the common view that

s lacks truth-value.

This view seems then to commit us to the consequence that *s* is neither true nor false:

_(*s* is true) & _(*s* is false).

But the first conjunct of this claim seems to be of the form:

s is not true,

and this is the liar sentence itself! To maintain the common view, you seem to be committed to using *and endorsing* the pathological sentence itself.

What is needed is to be able to have among your counterassumptions the view that the liar sentence has a truth value without having among your assumptions the negative claim that the liar does not have a truth value.

In a classical setting, the way in which you reject a view is to assert its negation -- but this does not work in the case of argumentation involving the paradoxes. In order to handle both assumptions and counterassumptions in arguments we probably need to distinguish two modes of assertion, along the lines of Woodruff, Peter, "Logic and Truth Value Gaps," in Karel Lambert (ed) *Philosophical Problems in Logic*, 1970, 121-42, and/or to distinguish assertion from denial, as in my "Assertion, Denial and the Liar Paradox," *Journal of Philosophical Logic* 13 (1984) 137-152. But such details go beyond the present paper.