Thinking about Thinking:

A Brief Introduction to Logic

Richard Oxenberg

I. Critical Thinking

Our course this semester will call upon us to apply our critical reasoning to pressing ethical and political questions of our day. The English word 'critical' is originally derived from the Greek word '*Kritikos*,' which means to discern or decide. Critical thinking, then, is a process of making sound decisions or judgments about the matter thought about. Not every critical examination involves matters of truth. An art critic might be more concerned to judge the beauty, or more generally, artistic merit, of the art piece under examination than its 'truth' per se. In philosophy, however, we are mostly concerned with truth claims. We can define critical thinking in a philosophical context, then, as the practice of trying to make sound judgments about claims concerning what is true and false.

II. Truth Claims

In philosophy, statements that make truth claims are called 'propositions,' insofar as they 'propose' something to be taken as true. It is the task of the critical thinker to try and determine whether such truth claims are warranted, that is, whether we should indeed accept their claims as true.

There are two ways in general to determine the truth of a truth claim. We might attempt to have a direct experience with the reality referred to in the statement. If someone says 'It is raining out,' for instance, we might go outside and look for ourselves, and on that basis decide whether or not the statement is true. Or we might judge the truth of a statement on the basis of its relationship to other statements we take to be true. For instance, someone might say, 'There are water droplets on the window, therefore, it must be raining out.' In this case, the truth claim 'it must be raining out' is supported, not by direct experience, but by the statement's relationship to another truth claim that, presumably, we can directly verify through experience. The statement 'There are water droplets on the window' is said to *imply* or *suggest* the truth of the statement 'It must be raining out.' Of course, we might agree or disagree that the latter statement is suggested by the former. Perhaps something other than rain caused the water droplets on the window.

III. Arguments

A logical argument is a series of statements that conclude with a truth claim, in which the truth of the concluding statement is said to be established on the basis of its relationship to the

other statements of the argument (as opposed to on the basis of a direct appeal to experience). The example given above ('There are water droplets on the window, therefore, it must be raining out') is a very simple example of a logical argument.

Every logical argument is made up of a *conclusion* (the statement whose truth the argument claims to establish, in this case, 'It must be raining out') and one or more *premises* (the statement or statements on the basis of which the truth of the conclusion is said to be established, in this case 'There are water droplets on the window').

In general, there are two types of logical argument. A *deductive* argument claims that its conclusion follows *necessarily* from its premises, such that it cannot be the case that the premises are true and the conclusion false. An *inductive* argument makes a more modest claim; it claims only that its conclusion is *more likely* to follow from its premises than not. The above argument would be a good example of an inductive argument. It would fail as a deductive argument because water droplets on a window do not necessarily imply rain. In what follows we will focus on the logic of deductive arguments and say more about inductive arguments later on.

IV. Logic, Validity, and Soundness

The word 'logic' is derived from the Greek word *logos*, which refers, in general, to language in its rational structure. Philosophical logic is a discipline that examines the rules and procedures for judging the quality of logical arguments.

There are, in general, two judgments that might be made concerning a deductive argument. We might judge the argument to be *valid* or *invalid*. And we might judge the argument to be *sound* or *unsound*. A *valid* argument is an argument whose conclusion follows necessarily from its premises. A *sound* argument is a valid argument whose premises are true, and whose conclusion, therefore, *must* be true. Note carefully that the fact that a deductive argument is valid does not in itself imply that the conclusion is true. A conclusion may validly follow from premises that are themselves false. In that case the argument would be *valid*, but *unsound*.

For example, consider the following argument:

P1 (premise 1): All cats wear tuxedos.
P2 (premise 2): All those who wear tuxedos play the piano.
C (conclusion): Therefore, all cats play the piano.

We know, of course (on the basis of direct experience), that neither the premises nor the conclusion of this argument are true. An examination of the logic of the argument, however, indicates that the argument is perfectly *valid*, that is to say, the conclusion follows necessarily from the premises such that, *were* the premises true, the conclusion would have to be true as well. We would say of this argument that, although its logic is *valid*, the argument itself is not

sound. For an argument to be sound it must be valid *and* all its premises must be true. Only sound arguments actually establish the truth of their conclusions.

Note also that an argument can be made up of nothing but true statements, and still be *invalid*. Consider the following argument:

P1: All cats are felines
P2: All felines are mammals
C: Therefore, all lions are mammals

All the statements in this argument are true, but the argument itself is not *valid*, because the conclusion is not implied by the premises. In fact, the premises don't make any mention of lions at all. Our knowledge of the truth of this conclusion, then, cannot be based on our knowledge of the truth of the premises, which is what is required for an argument to be valid.

This means that there are two ways, in general, for us to challenge the soundness of a deductive argument. We can deny that the argument is valid, and thus deny that the conclusion follows from the premises. Or we can deny that one or more of the premises are true. In either case, the argument would not be sound, that is, it would fail to establish the truth of its conclusion.

V. Categorical Logic

The discipline of logic is often divided, very broadly, into two categories: Categorical and Propositional logic. Categorical logic, which is the more elementary, deals with arguments made about the relationship between *categories* or *classes* of things. In language, every noun that is not a proper noun refers to a class or category of thing. For instance, the noun 'chair' (without a modifier) refers, not to any particular chair, but to the class or category of chairs. Categorical logic works through indicating relationships of inclusion and/or exclusion among classes or categories of things. For instance, the class 'chair' is included in the class 'furniture' (i.e., chairs are kinds of furniture). From this we know that anything that is true of all things included in the class 'furniture' will also be true of all things included in the class 'chair.'

The simplest form of categorical argument is the *categorical syllogism*. In its broadest use, the word 'syllogism' can refer to any deductive argument. A *categorical* syllogism, however, more narrowly refers to an argument that relates two class *terms* by showing their relation to a third class term (called the 'middle term').

An example of a categorical syllogism would be as follows:

P1: All human beings are mortal
P2: All Greeks are human beings
C: Therefore, all Greeks are mortal.

In this argument we have three terms ('Greeks,' 'mortals,' and 'human beings'), each of which refers to a different category or class of thing. The argument's conclusion claims that all entities that fall within the category referred to by the term 'Greeks' also fall within the category referred to by the term 'mortal.' It makes this claim on the basis of a relationship claimed between each of these classes and a third class (named in the 'middle term') which does not appear in the conclusion, but appears in each of the premises, the class 'human beings.' Given that all entities that fall within the category 'human beings' fall within the category 'mortal' (premise 1), and given that all entities that fall within the category 'Greeks' fall within the category 'human beings' (premise 2), it follows necessarily that all entities that fall within the category 'Greeks' also fall within the category 'mortals.' This is a valid argument.

On the other hand, consider the following:

P1: Some human beings are mortal.
P2: All Greeks are human beings.
C: Therefore, all Greeks are mortal.

The argument, modified in this way, is no longer valid. The fact that some human beings are mortal and all Greeks are human beings does not in itself imply that all Greeks are mortal, since, on the basis of the premises alone, it remains possible that *some* human beings are *not* mortal, and that some of these are Greeks. So this is an invalid argument. Notice, however, that all the statements of the argument are *true*. We would say of this argument that, although all its statements are true, it is nevertheless not *valid*, because the truth of its conclusion cannot be established on the basis of the truth of its premises.

VI. Propositional Logic

Whereas the premises of categorical syllogisms assert a relationship between categories of things, the premises of propositional arguments assert relationships between propositional truth claims themselves, arguing that one proposition must be true on the basis of a relationship of implication that exists between that proposition and another. There are two major forms of propositional syllogism: the hypothetical and the disjunctive syllogism.

VI.1 The Hypothetical Syllogism

A hypothetical *statement* asserts a necessary relationship between the truth of two propositions, expressed in the form of an 'If…then…' statement. The first premise of every hypothetical syllogism is always a hypothetical statement. The second premise asserts or denies the truth of one of the propositions of the hypothetical statement. Finally, the conclusion asserts or denies the truth of the other proposition of the hypothetical statement on the basis of the premises.

The two major types of hypothetical syllogism are called '*modus ponens*' (meaning 'the way of affirmation') and '*modus tollens*' ('the way of negation').

In the *modus ponens* syllogism the first premise (the hypothetical statement) asserts that one proposition (called the 'antecedent,' given in the *if* clause) necessarily implies another (called the 'consequent,' given in the *then* clause), the second premise affirms that the antecedent is true, and, on the basis of these premises, the conclusion asserts that the consequent is true.

Here is an example of a simple *modus ponens* argument:

P1: If the ground is wet, then it rained last night.
P2: The ground is wet.
C: Therefore, it rained last night.

P1 is the hypothetical statement, with 'the ground is wet' as its antecedent and 'it rained last night' as its consequent. P2 affirms that the antecedent of the hypothetical statement in P1 is true. Therefore, the conclusion asserts, the consequent of the hypothetical statement must be true as well.

In the *modus tollens* form of the hypothetical syllogism the second premise asserts that the consequent of the first premise is *not* true, and on this basis concludes that the antecedent must not be true as well.

A simple example of a *modus tollens* argument would be as follows:

P1: If the ground is wet, then it rained last night.
P2: It did not rain last night.
C: Therefore, the ground is not wet.

Again, this is a valid argument. *Given* that the consequent of a hypothetical statement is necessarily implied by its antecedent, it necessarily follows that if the consequent is *not* true then the antecedent must *not* be true as well.

VI.2 The Disjunctive Syllogism

The second major form of propositional argument is the disjunctive syllogism. In the disjunctive syllogism the first premise (the disjunctive statement) asserts that one *or* another proposition must be true. The second premise asserts that one of the propositions of the disjunctive statement is not true. On the basis of this the conclusion asserts that the other proposition of the disjunctive statement must be true.

For example:

P1: Either it rained last night *or* the Red Sox played last night.
P2: It did not rain last night.
C: Therefore, the Red Sox played last night.

This is a valid argument, although, again, the truth of the premises might be questioned. For instance, if it were not baseball season, or if the Red Sox were not scheduled to play last night, it might be the case that both propositions of the disjunctive statement are false, which would make premise 1 false. Still, the argument is valid insofar as the conclusion follows necessarily from the premises.

VII. Inductive Reasoning

Inductive arguments make weaker claims than deductive arguments. Whereas a deductive argument claims that the truth of its conclusion *necessarily* follows from the truth of its premises, an inductive argument merely claims that its conclusion is *more likely* to be true than not, given the truth of its premises. Given this, deductive validity is not the criterion by which the merit of an inductive argument can be judged. A good inductive argument, thus, is not spoken of as 'valid' but as 'cogent' (i.e., convincing). Nevertheless, it is still possible to employ logical principles in evaluating inductive arguments.

An example of an inductive argument might be:

P1: Most human beings have two legs.
P2: All Greeks are human beings.
C: Therefore (it is probably the case that), most Greeks have two legs.

Notice that this would fail as a deductive argument. It might be the case, on the basis of the premises alone, that the majority of Greeks are among the minority of human beings who do not have two legs. Nevertheless, barring any good reason to think that the population of Greeks differs from the general population of human beings in this respect, we might well judge this argument to be convincing. The conclusion doesn't *necessarily* follow from the premises, but the premises provide good support for the conclusion.

Notice that if we add another premise to the inductive argument above we can turn it into a valid deductive argument:

P1: Most human beings have two legs.
P2: All Greeks are human beings.
P3: The proportion of Greeks who have two legs is equivalent to the proportion of human beings in general who have two legs.
C: Therefore, most Greeks have two legs

Now the argument is deductively valid. One powerful way of assessing the cogency of inductive arguments is to try to determine what *additional* premises would be needed to turn them into valid deductive arguments, and then make a judgment as to what extent it is reasonable to regard these additional premises as true. The cogency of an inductive argument will be proportional to the probability that the additional premises required to turn it into a valid deductive argument are true.

VIII. Ordinary Language

Most argumentation in real-world discourse is not presented in the strict logical forms we have been reviewing above. Such logical forms are studied to aid us in assessing the logical structure of real-world arguments, but, to do so, we often have to translate from the arguments as given in ordinary language to the forms that reveal their logical structure. The value of doing this is that, in examining the logical form of the arguments, we are in a better position to assess their soundness, and, thus, make a sound judgment as to the truth of their conclusions.

Let's take a simple example from one of the arguments pertaining to the abortion controversy. One of the most popular arguments against abortion might be stated in ordinary language as follows:

'Abortion is murder pure and simple! Everyone believes that killing innocent human beings is wrong, and fetuses are certainly innocent. So to kill a fetus is obviously wrong.'

This argument can be expressed in the form of a simple categorical syllogism:

P1: All innocent human beings are beings whom it is wrong to deliberately kill.
P2: All fetuses are innocent human beings.
C: Therefore, all fetuses are beings whom it is wrong to deliberately kill.

Notice that to exhibit this argument's logical form we have to make explicit what is implicit in the argument as given in ordinary language: that fetuses are human. When we do, however, we can see that the argument is logically *valid*, that is, the conclusion follows necessarily from the premises.

Anyone wishing to deny the conclusion of this argument, then, must challenge one or another of its premises. Most who would defend the morality of abortion would challenge the truth of premise 2, and deny that fetuses are 'innocent human beings.' We will read an article later in this semester that will challenge this argument by denying premise 1. The point here, however, is that by revealing the logical structure of the argument we put ourselves in a better position to know how to critically examine it.

IX. Informal Fallacies

The examination of the logical structure of ordinary language arguments is often made difficult by what are known as 'informal fallacies.' Informal fallacies are modes of argumentation, found in ordinary language arguments, that often have persuasive power but do not have logical legitimacy. Their persuasive power is generally a function of logical confusion, linguistic ambiguity, or appeal to irrelevant emotions and/or prejudices. In order to assess the logic of ordinary language arguments we often, first, have to identify and eliminate any informal fallacies they may contain.

Logicians have named and described some of the most common of these informal fallacies in order to help us avoid committing them ourselves, and in order to help us identify them in arguments we are assessing. In the following I list and describe some of the most common informal fallacies and provide examples culled from popular arguments we hear with some frequency in ordinary language debate.

IX.1 Begging the Question

An argument is said to 'beg the question' when its conclusion is merely a restatement of one if its premises. The purpose of an argument, of course, is to establish the truth of a conclusion by showing its relationship to premises whose truth may be less controversial, or simply more well established, than the truth of the conclusion itself. To derive a conclusion from *itself*, then, defeats the purpose of the argument and proves nothing.

An example of an argument that 'begs the question' might be as follows:

> 'We can be sure of the truth of the Bible because God is its author, and God would never lie. Further, we can be sure that God is the author of the Bible because the Bible says so itself! Therefore, the truth of the Bible is beyond question.'

This argument seeks to establish the truth of the Bible on the grounds that God is its author, but then seeks to establish that God is the author of the Bible on the grounds of (the truth of) the Bible. Of course, if the Bible's truth is questionable we can also question its claim to be authored by God. This is a classic example of 'begging the question,' since the conclusion of the argument, that the Bible is true, depends upon our accepting this very conclusion as a premise of the argument.

IX.2 Ambiguity

Some ordinary language arguments are fallacious because they employ terms that are ambiguous, or equivocal, i.e., that have more than one meaning depending upon context, and then use this ambiguity to make inferences that would not be valid were the ambiguity resolved.

An example of this can be seen in a very common argument against Affirmative Action programs:

> 'We know that discrimination on the basis of race is unjust. Affirmative Action programs discriminate on the basis of race. Therefore, Affirmative Action programs are unjust.'

This *seems* like a sound argument, and can easily be translated into the terms of a simple, valid, categorical syllogism. Nevertheless, it commits the fallacy of ambiguity given that 'discrimination on the basis of race,' which indicates racism in certain contexts, does not indicate it in others. For instance, a doctor who gives different medical tests to African-Americans than to Caucasians, on the grounds that one group is more susceptible to a certain disease, would be discriminating on the basis of race, but we would not consider this to be unjust. Thus, it is only

discrimination on the basis of race for the sake of *persecution* that is unjust. Once we make this clear the argument's premises prove no longer true, or the argument itself proves no longer valid (depending upon how we reframe it), given that Affirmative Action programs do not discriminate for the sake of persecution. Of course, this does not tell us that Affirmative Action programs discriminate justly, only that this particular argument against them is flawed.

IX.3 Ad Hominen

'Ad Hominem' is Latin for 'to the man.' An argument commits the 'ad hominem' fallacy when it argues against an opponent's position on the basis of some character flaw of the person holding that position, rather than on the basis of the logic of the opponent's argument itself.

Example:

> 'Libertarians claim that an unregulated free market makes for a more just society, but the reason they claim this is because they are greedy and don't want their greed curbed through regulation, so their claims can be dismissed.'

This is an example of a fallacious 'ad hominem' argument. Maybe it is and maybe it is not the case that libertarians are greedy and/or that this motivates their ideology, but that is beside the point. The question of whether or not an unregulated free market makes for a more just society must be decided on the basis of some conception of justice and some examination of the way justice is affected by the marketplace, not on the basis of the character traits, or even the motives, of those who defend free markets.

IX.4 Appeal to Ignorance

The fallacy of 'appeal to ignorance' argues from the *lack* of evidence or knowledge of something to the fact that it is not, or is not likely to be, true. In fact, a lack of evidence proves nothing at all (that is why we call it a 'lack'), and, in and of itself, cannot support a truth claim one way or another.

Example:

> 'There has been no evidence whatsoever that electronic voting has been used for vote tampering, therefore we can conclude that it has not been so used."

This is a fallacious argument based upon an appeal to ignorance. The lack of evidence for vote manipulation is not, in and of itself, evidence that the vote has not been manipulated, unless one can make and support the *further* claim that if there *had been* vote tampering, there *would be* evidence of it. But the truth of this latter claim is not implied by a mere lack of evidence. Perhaps evidence was not properly sought. Or perhaps (as opponents of electronic voting indeed claim) the nature of electronic voting is such that it makes it difficult, or impossible, to gather the necessary evidence to determine the matter one way or the other. In any case, *lack* of evidence is never, in and of itself, *evidence*.

IX.5 Tu Quoque

'Tu Quoque' is Latin for 'you too.' We might think of this as the 'two wrongs make a right' fallacy. The fallacy of 'tu quoque' involves suggesting that your argument, position, or behavior is correct or acceptable on the grounds that it is no worse than your opponent's argument, position or behavior.

Example:

> 'I think it is perfectly acceptable for me to lie on occasion, despite what Kant may say. Indeed, Kant's claim that it is always wrong to lie is proven false by the fact that Kant himself was known to lie on more than one occasion.'

The fact that Kant himself may have lied on occasion has no logical bearing on the question of whether his arguments concerning the morality of lying are sound. Maybe Kant was a hypocrite. This in itself would not make his moral theory unsound. (By the way, I am only using this as an example. I have no idea whether Kant was ever known to lie!).

IX.6 Appeal to Authority or Tradition

In this fallacy one attempts to establish the truth of a position by appeal to a reputed, but questionable, authority or tradition. Note that it is not always illegitimate to appeal to an authority when there is good reason to believe that the authority is a reliable witness to the truth of the question at issue. In that case, though, the appeal is not to the authority *per se*, but to the authority's *knowledge* of the issue at hand.

Example:

> 'Homosexual activity is morally wrong, and we know this because there are any number of passages in the Bible that condemn it.'

The cogency of this argument depends entirely upon the reliability of the Bible as an infallible witness to what is right and wrong. Given, however, that the argument itself provides no reasons to believe that the Bible can always be relied upon in this way, it commits the fallacy of 'appeal to authority.' To make the argument work, some further argument concerning the reliability of the Bible is needed.

IX.7 Hasty Generalization

In the fallacy of 'hasty generalization' one argues from one, or a few, particular instances of something to a general rule without providing sufficient reason to believe that the general rule is indicated by the instances.

Example:

> 'Evangelical Christians are a bunch of charlatans and hypocrites. For proof of this just look at Jimmy Swaggart, who visited prostitutes all the time he was preaching, or James Baker, who bilked his followers out of millions of dollars!'

Of course, the fact that there may be some dishonest or hypocritical evangelical Christians is insufficient reason to assume that all evangelicals are hypocrites. To avoid the fallacy of hasty generalization one needs to provide some reason to think that what is true in one or a few instances is an *essential* feature of the class in question.

IX.8 False Dilemma

This is a fallacy that results from constructing an untrue disjunctive statement. Remember that a disjunctive statement claims that one or another proposition *must* be true. But, of course, this is not always the case. Sometimes there are alternatives to the propositions of the disjunction that would allow both to be false, in which case the negation of the one would not imply the truth of the other.

Example:

> 'Either there were weapons of mass destruction in Iraq or G.W. Bush lied to us. There were no weapons of mass destruction in Iraq, so Bush must be a liar.'

As we see, this argument takes the form of a simple disjunctive syllogism. But the disjunction itself might be challenged by suggesting that Bush was honestly mistaken about whether or not there were weapons of mass destruction in Iraq, in which case it might be true *both* that Bush did not lie *and* that there were no WMD in Iraq.

We might eliminate the false dilemma by reframing the disjunctive syllogism as follows: 'Either Bush was ignorant that there was good reason to question the existence of WMD in Iraq, or Bush lied to us. Bush was not ignorant about this. Therefore, Bush lied to us.' This argument no longer commits the fallacy of false dilemma, but note how the evidentiary criteria for establishing its soundness has shifted. We now need to establish something about what Bush did or didn't know. This is a good example of how clarifying our arguments can help get us on the right track in trying to determine the truth.

IX.9 Naturalist Fallacy

The naturalist fallacy is of particular relevance to moral arguments. It argues that something *ought* to be the case, morally, on the grounds that something *is* the case, in fact. That this is a fallacy becomes obvious upon reflection on the meaning of moral claims. Moral claims, as Plato was perhaps the first to make clear, cannot be derived from what is factually the case. For instance, the fact that some people *are* murdered certainly doesn't imply that they *ought* to be.

Example:

> 'The capitalist system is based on the principle of survival of the fittest. This is the same principle that natural evolution is based upon. The capitalist system, thus, can be seen as simply an outgrowth of the system of natural evolution itself, and, therefore, cannot be regarded as unjust.'

There are many questionable things about this argument, but, in particular, it commits the naturalist fallacy. Even if we allow that capitalism is an outgrowth of nature (a questionable claim in itself), the mere fact that something is a fact of nature does not, in and of itself, imply that it is morally right or good. Would we say, for instance, that because people naturally get sick it is 'right' that they get sick? Claims about right and wrong (or justice and injustice) must be based upon some theory or conception of morality, not upon what is or is not observed to be the case in nature or elsewhere.

IX.10 Appeal to the Masses

In this fallacy an argument is made from the popularity of a belief to its truth. Of course, the mere fact that a belief may be popular does not, in and of itself, guarantee its truth.

Example:

> 'In this day and age almost no one believes any longer that premarital sex is wrong. Therefore, claims that it is wrong need not be taken seriously.'

The popularity of a belief is determined by many factors that might have nothing at all to do with the truth of what is believed. Thus, it is never legitimate to argue from the mere popularity of a view to its truth.

IX.11 Genetic Fallacy

An argument commits the genetic fallacy when it attempts to refute, or affirm, an argument on the basis of that argument's origin (or 'genesis') rather than on the basis of the merit of the argument itself.

Example:

> 'Thomas Aquinas' arguments for the existence of God may have seemed convincing in his day but we are no longer living in the Middle Ages, so we can dismiss them.'

The mere fact that Aquinas wrote in the Middle Ages is not a legitimate reason to dismiss his work. Even if one wishes to assert that there are flaws in Aquinas' arguments that are typical of

Medieval thinkers in general, one would still need to show what those flaws are. It is fallacious to attempt to refute (or affirm) an argument on the basis of its genesis alone.

X. Conclusion

Real-world decisions, sometimes of enormous gravity (such as whether or not we go to war, or whether or not we subject a person to imprisonment or death) must be made on the basis of rational argumentation, as the truth claims pertaining to such decisions are rarely verifiable through simple empirical observation. It is incumbent upon us as decision-makers (and all citizens of a democracy are decision-makers) to strive to understand and evaluate the arguments we hear from others, and those we frame ourselves, in a way that will enable us to make decisions soundly. As citizens of a democracy, and simply as human beings, we are called upon to make judgments concerning the most weighty moral and political matters. If we are to do so responsibly, it is imperative that we learn to think carefully and well about the issues that confront us.