

# How to Read Philosophically

The aims of reading philosophically .....	1
What is the issue? .....	2
What is the conclusion?.....	2
What are the premises?.....	3
How, by the author’s lights, do the premises support the conclusion? .....	4
Independent versus dependent premises .....	5
Intermediate and ultimate conclusions .....	7
Some pointers on paraphrasing .....	8
Do the premises adequately support the conclusion?.....	9
Induction, deduction, and counterexamples .....	11
Are the premises true?.....	12
Summary of rules .....	14

## The aims of reading philosophically

We all know how to read, but not all of us know how to read *effectively*. Furthermore, just as there are specialized forms of writing for certain disciplines and fields, there are certain forms of *reading* that are equally specialized. In what follows, I offer some strategies for *philosophical* reading. From my own opinionated perspective, these strategies are effective in fields far beyond philosophy; indeed they are useful for reading most nonfiction.

Obviously, with any reading, you want to understand what you’ve just read. As we all know, there is a difference between “merely memorizing” what you’ve read and “genuinely understanding” it. However, you’ll find very different answers (many of them vague and unenlightening) if you push most people on what they mean by genuine understanding. I want to stress that understanding is not a subjective feeling of “Eureka!” or having a “light go off in your head.” Understanding is a very public, objective thing. Fundamentally, a person only understands what he/she has read if she can answer all of the following questions:

- a. *What is the issue?*
- b. *What is the conclusion or thesis?*
- c. *What are the premises?*
- d. *How, by the author’s lights, do the premises support the conclusion?*
- e. *Do the premises adequately support the conclusion?*
- f. *Are the premises true?*

People can have Eureka moments and be incapable of answering these questions, and others might not have Eureka moments at all, but are perfectly capable of answering these questions. Which person would you want to study with? We all prefer the latter person, because she has useful information; the first person just feels good about himself, but that won’t help us pass a test, engage in fruitful discussions, or write a good paper. So the first of many boldface strategies is this: **when taking notes on a text, answer these six questions.** The remaining strategies and discussions are merely pointers on how to realize this goal.

## What is the issue?

An issue is the main question that the author is asking. Almost all of us can identify an issue by skimming an article, and it's often a good idea to **do a quick skim to identify the issue before doing a closer reading.**

But how do you go about looking for an issue? One thing to note is that **good writers will almost always introduce the issue in the first few sentences of their works.** Typically, a good introduction “motivates a problem,” i.e., it tries to show what the issue is. Closely related, **issues are often things of which the author takes pains to underscore their importance.** For example, if an author says “Terrorism is the biggest challenge for global harmony in the 21<sup>st</sup> century,” there's a good chance that terrorism will figure in the issue. Similarly, **issues are often things that the author treats as in tension or problematic in some sense.** For example, if the author writes, “Balancing issues of national security with civil liberties is becoming increasingly complicated in a post-9/11 world,” national security, civil liberties, and terrorism are likely to figure in the issue.

But once you've identified key concepts that will figure in the issue, your work is not yet done. **Phrase your issue as a question, not as a word or a phrase.** Words and phrases typically can play more roles than questions can. For example, if you say that a passage is tacking the issue of “justice,” any of the following might be fair interpretations “Is a particular policy or practice just?” “What is the definition of justice?” “Why is justice better than injustice?” etc. Intuitively, these are *distinct* issues. Most of us probably haven't made an issue out of defining justice or justifying its superiority to injustice, but many of us have considered whether certain policies or practices are unjust (e.g., slavery, wiretapping, etc.). Using a word/phrase as an issue runs the risk of conflating different questions, some of which aren't relevant to the issue that the author is considering. Sometimes authors aren't clear about which issues they're concerned with too, so **look out for when authors change the issue.**

Of course, there are better and worse ways to formulate a question, so **make your question as precise as possible.** Suppose that we are reading a text, and we have a hunch that the issue amounts to: “Is wiretapping just?” Well, certainly some kinds of wire-tapping are unjust, but these are typically assumed as such by the author. For example, if a stalker wiretaps a celebrity for no other reason than his perverse obsession, everyone would agree that this is unjust. So try to frame your question more explicitly, e.g., “Is wiretapping suspected terrorists' phones just?” Perhaps you can get even more specific than this, e.g., “Is wire-tapping people who have unknowingly donated money to Al-Qaeda just?” The more specific the question, the easier it is to see which parts of the passage are relevant and which are serving some other role. Closely related to this is another strategy: **wherever possible, frame the issue as a yes-no question.** Ultimately, a precisely framed issue should ask us to evaluate whether a particular position is correct or incorrect, i.e., to either say Yes or No to the proposal being considered. Not all passages admit this. Sometimes the very point of a passage is to show how the existing positions draw too strict of a dichotomy, and this is important to remember as well. But, *if* you can frame the issue as a Yes-No question, do it!

## What is the conclusion?

If you've formulated the issue as a precise question, the conclusion is simply the *answer* to that question. More importantly, as we noted above, the conclusion is also the *thesis* of the passage. While the order suggests that you identify the issue first and then formulate the conclusion in response, in practice, you should **go back and forth between formulating**

**tentative guesses about the issue and formulating the conclusion, until you have the most precise formulation of both.** Typically this process involves generating good guesses about either the issue or the topic in response to various textual cues.

How do you go about identifying the conclusion? Well, very often they will have special words associated that flag them, so you should **always look for conclusion-indicators, i.e., words such as “therefore,” “thus,” “hence,” “so,” “as a result,” “it follows that,” etc.** For example, if after a long discussion, an author writes “So it would seem that wiretapping is not the cure-all for our national security woes,” you can bet that the long diatribe preceding it culminates in this point. However, **not all conclusions are indicated as such.** Sometimes, context makes conclusion-indicators unnecessary, since it is clear that the conclusion follows. Furthermore, not all conclusion-indicators flag the *main* conclusion of the text, so **be careful not to mistake intermediate conclusions, i.e., statements that both are justified by other reasons and help to justify the main or ultimate conclusion, for ultimate conclusions.** These will often have conclusion-indicators but are not going to help you to see the big picture. We’ll return to the important positive role that intermediate conclusions can play when we talk about how the premises support the conclusion.

## What are the premises?

Once you’ve identified the issue and conclusion, you’re now ready to ascertain what the author’s premises, *evidence*, or *reasons* for believing that conclusion. Note that you should **ascertain an author’s reasons before you criticize them.** It’s a much more significant intellectual accomplishment (and shows greater understanding) if you refute an author’s ideas after showing that you’ve faithfully reconstructed his/her reasoning than if you caricature his ideas, what is typically called a straw man argument. **Avoid straw man arguments.** Closely related to this, you should **be charitable**, trying to interpret the passage in a manner that makes the premises as plausible as the text will allow.

While there is some give and take between formulating issues and formulating conclusions, it’s almost never a good idea to identify the premises before the conclusion. The reason for this is that an argument by definition only has one (ultimate) conclusion, while it may have an infinite number of premises, so in order to avoid losing the forest for the trees, you should **always identify the conclusion before you identify the premises.** However, like conclusions, you should **look for premise-indicators, e.g., “since,” “because,” “it follows from,” “as supported by,” etc.** As with conclusions, **not all premises will be indicated as such.**

It’s tempting to think that everything that isn’t the conclusion is a premise. Resist this urge; even very good writers go on (often very illuminating) tangents, engage in stylistic flourishes, provide background information, and generally include information that is not immediately relevant to establishing their conclusions. Furthermore, **distinguish details from what is absolutely essential to supporting the conclusion.** For example, if someone writes that “Wiretapping dates back to the earliest telegraphs” but her conclusion is that “Wiretapping people who unwittingly gave funds to Al-Qaeda sympathizers is unjust,” it’s clear to see that this piece of history does not support the conclusion. **To ascertain what is necessary, engage in critical imaginative steps, e.g.,** suppose that wiretapping *didn’t* date back to the telegraph, and ask yourself whether this would make it any more or less just in the contemporary context. Quite clearly, it doesn’t, since it could have been equally just or unjust in the past.

Closely related, there's a certain art to **stating premises as generally as possible without being too vague**. Specifically, you want premises that are general enough not to get bogged down in irrelevant details but are precise enough to support the conclusion, which as an answer to a precisely formulated issue, will also be quite precise. Thus, **don't lose sight of the conclusion when identifying the author's premises**.

## How, by the author's lights, do the premises support the conclusion?

Once you've identified what the premises are, you have to ascertain how the author *uses* them to support the conclusion. Here, the key issue is to **get familiar with different patterns of argument**, i.e., typical ways in which premises support conclusions. Here is a list of common argument patterns:

A. Inductive Arguments			
Name of Argument	Form(s)	Defeater	Rules
Argument by example	Some (representative) $F$ 's are $G$ 's. $\therefore$ (Probably) Most $F$ 's are $G$ 's.	Sample not Representative	Give more than one example; use representative examples; note background info; consider counterexamples
	All observed $F$ 's are $G$ 's! $\therefore$ (Probably) All $F$ 's are $G$ 's.		
	This (random) $F$ is a $G$ . $\therefore$ (Probably) Most $F$ 's are $G$ 's.		
Argument by analogy	$a$ is an $X$ . Both $a$ and $b$ are similar with respect to $F, G, H, \dots$ $\therefore$ So $b$ is (probably) an $X$ .	Irrelevant analogy	Be precise about points of similarity, target, and points of difference
Argument from authority	Some source asserts or indicates that $p$ . $\therefore$ (Probably) $p$ .	Unreliable source	Cite sources; seek and cross-check informed and impartial sources; avoid personal attacks
Inference to the Best Explanation	$E$ $H$ best explains $E$ . $\therefore$ (Probably) $H$ .	Alternative hypothesis	Look for evidence that falsifies the evidence; imagine other explanations
Causal arguments	(Under appropriate, controlled conditions,) there is a strong correlation between $X$ and $Y$ . $X$ precedes $Y$ in time. $\therefore$ $X$ (probably) causes $Y$ .	Spurious correlation	Explain how cause yields effect; propose likeliest cause; correlation $\neq$ causation; search for common causes; watch for causal direction; look out for complex causes

<b>B. Deductive arguments</b>		
Modus ponens	If $p$ then $q$ . $p$ . $\therefore q$ .	<b>Not applicable</b>
Modus tollens	If $p$ then $q$ . Not $q$ . $\therefore$ Not $p$ .	
Hypothetical Syllogism	If $p$ then $q$ . If $q$ then $r$ . $\therefore$ If $p$ then $r$ .	
Disjunctive Syllogism	Either $p$ or $q$ . Not $p$ $\therefore q$ .	
Dilemma	Either $p$ or $q$ . If $p$ then $r$ . If $q$ then $s$ . $\therefore r$ or $s$ .	
Categorical Syllogism	All $F$ 's are $G$ 's. $a$ is $F$ . $\therefore a$ is $G$ .	


We'll talk a bit more about inductive and deductive reasoning and the right two columns below. But the key point is that once you get familiar with these forms, start **paraphrasing parts of passages as instances of these forms of argument**, as almost all passages can be turned into complexes of these simpler forms of argument. But what does it mean to say they are *complexes*? We can point to two crucial ideas: independent versus dependent premises and intermediate versus ultimate conclusions.

### Independent versus dependent premises

Sometimes, there is one, tightly interconnected argument for a conclusion, in which any individual premise by itself would not provide a good reason to accept the conclusion, but when that premise is combined with all of the other premises in the argument, the combination is quite compelling. We call these *dependent premises*. For example:

- (1) Anyone under 18 is not allowed on the premises.
- (2) Sally is under 18.
- (3)  $\therefore$  Sally is not allowed on the premises.

Now if it were the case that (1) is true and (2) is false, i.e., anyone under 18 is not allowed on the premises but Sally is 18 or older, then this would be a bad argument. Similarly, if (1) were false and (2) were true, i.e., people under 18 *are* allowed on the premises and Sally is under 18, then this would also be a bad argument. Thus, (1) and (2) *depend* on each other if they are to support (3). Pictorially, we represent dependent premises in the following manner:



Contrast this with the following.


ARGUMENT A:

- (1) Columbia is an elite institution of higher learning.
- (2) Elite institutions of higher learning should pursue the truth.
- (3) If someone is a Holocaust denier, then he is opposed to the pursuit of truth.
- (4) Ahmedinejad is a Holocaust denier.
- (5)  $\therefore$  Columbia should not invite Ahmedinejad to speak.


ARGUMENT B:

- (1) Columbia is an elite institution of higher learning.
- (6) Elite institutions of higher learning should promote international peace.
- (7) If someone threatens nuclear attacks on the U.S. and its allies, then he does not promote international peace.
- (8) Ahmedinejad threatened nuclear attacks on the U.S. and its allies.
- (5)  $\therefore$  Columbia should not invite Ahmedinejad to speak.

Premises (2)-(4) and (6)-(8) provide *independent reasons* for accepting (5) as a conclusion, since regardless of whether or not Columbia is committed to promoting international peace, if it is committed to pursuing truth, ARGUMENT A would be an argument for (5). Analogously, if Columbia is not committed to pursuing truth, ARGUMENT B would still be an argument for (5). We represent independent premises in the following manner:



For the Columbia argument, we diagram it as follows:



Here, the left cluster of number propositions corresponds to ARGUMENT A while the right corresponds to ARGUMENT B. The big lesson here is to **where appropriate, paraphrase arguments such that independent premises are carefully distinguished from dependent premises**. This becomes especially important when we turn to *evaluating* arguments, since showing one dependent premise to be false is sufficient to undermine a whole argument, but showing one independent premise to be false is not.

### ***Intermediate and ultimate conclusions***

Earlier, we suggested that there should be *no more than one* (ultimate) conclusion and *at least one* premise in any argument. What do we mean by an *ultimate* conclusion? Certain propositions may be conclusions to one argument but then serve as premises in a larger or subsequent argument, these are often called *intermediate conclusions*. It's important that you **recognize the larger argument of which they're a part or else you'll lose the forest for the trees**. For example, we might break apart ARGUMENT A above in the following manner:

ARGUMENT C:

- 1) All elite institutions of higher learning are committed to the pursuit of truth.
- 2) Columbia University is an elite institution of higher learning.
- ∴ 3) Columbia University is committed to the pursuit of truth.


ARGUMENT D:

- 4) If someone is a Holocaust denier, then he is not committed to the pursuit of truth.
- 5) Ahmedinejad is a Holocaust denier.
- ∴ 6) Ahmedinejad is not committed to the pursuit of truth.


ARGUMENT E:

- 7) If an institution is committed to the pursuit of truth, then it should not invite anyone not committed to the pursuit of truth to speak.
- 2) Columbia University is committed to the pursuit of truth.
- 6) Ahmedinejad is not committed to the pursuit of truth.
- ∴ 8) Columbia should not invite Ahmedinejad to speak.


Thus, (3) is the conclusion of ARGUMENT C, (6) is the conclusion of ARGUMENT D, and both are premises of ARGUMENT E. Therefore, (3) and (6) are *intermediate conclusions* and 8), the conclusion of ARGUMENT E, is the *ultimate conclusion*. A simple intermediate conclusion structure might look like this:



Where (2) is an intermediate, and (3) an ultimate conclusion. Our own monster, involving ARGUMENTS C, D, and E, looks like this.



In principle, you could insert intermediate conclusions nearly everywhere, but note that they're largely dispensable, i.e., if you have your first premises, you can get your ultimate conclusion. So, **unless the intermediate conclusion *really, really* helps you to understand the flow of the argument, leave it out.** It's not necessary. In principal, we could have diagrammed the argument as follows:



As with our discussion of "What are the premises?" the goal is to be as general as possible without being too vague, so if you can capture the same argument with fewer premises, kudos to you!

### ***Some pointers on paraphrasing***

It would be a bit much to ask you to actually *diagram* all of your arguments as we've done here. Primarily these serve as visual aids to help you grasp different structures of complex arguments. What you should be doing is *paraphrasing arguments*. In addition to the aforementioned strategies, here are some other rules of thumb to help you in this enterprise:

**A) List the premises in an *order* which makes the structure of the argument clear, minimally in *standard form*;**

- 1) Standard form is anything that follows the structure of the arguments listed in the table above.

**B) *Simplify* the language of the original text by trading out more elliptical and counterintuitive language for more concrete and concise language;**

**C) *Eliminate irrelevant propositions*;**

- 1) A proposition is irrelevant if it is neither a premise nor a conclusion of an argument, e.g., our example of the history of wiretapping above.

**D) Provide *uniformity* of terms and language;**

- 1) You're sometimes told in other classes not to be too repetitive with certain terms. However, there is a tradeoff: the more synonyms you use to keep your language 'varied,' the more likely you are to be imprecise with your language. Many synonyms carry slightly different connotations and meanings, and it is often hard to keep all of these in mind when you're writing. In philosophy, it's generally encouraged that the same terms are used throughout, and even then, very smart people will miss subtle equivocations in their use of a term. Furthermore, note that uniformity of language makes the argumentative structure more transparent. Consider the following:

- i. If murder is wrong, then eating meat is wrong.
- ii. Murder is morally abhorrent.



iii. So eating meat is a violation of the sanctity of life.

Is this a good argument? Maybe, maybe not, but note that the following is certainly a *clearer* argument and seems to say the same thing:

iv. If murder is wrong, then eating meat is wrong.

v. Murder is wrong.

vi. So eating meat is wrong.

All we did is insert “is wrong” for some of the more colorful variants in the original argument.

**E) Make *implicit premises* explicit, provided they are necessary to get an argument to fit one of the aforementioned forms;**

- 1) By far the most common version of this is inserting a conditional to get a *modus ponens*, e.g., suppose a passage said, “Murder is wrong. So eating meat is wrong.” We can turn this into a tighter argument by adding the premise “If murder is wrong, then eating meat is wrong.”

## Do the premises adequately support the conclusion?

At this point, you’ve given the author the benefit of the doubt, and tried your best to reconstruct his/her argument as clearly as possible by providing a precise formulation of the issue and conclusion he/she is advancing by using the most common and effective argument forms available to reconstruct her reasoning, distinguishing dependent from independent premises, and paraphrasing with a number of effective heuristics. You’re finally at a point to ascertain if the author has made a *good case* for his/her conclusion. The first way to do this is to examine if the premises *support* the conclusion, the second way, to be examined in the next section, is whether the premises are *plausible* in their own right. Note that perfectly plausible claims might not support a conclusion. For example,

If Tweety is a mammal, then Tweety is warm-blooded.

Tweety is warm-blooded.

So Tweety is a mammal.

The premises are true, but they do not support the conclusion, because Tweety is a bird, and birds are also warm-blooded. What this points to is that fully understanding a text involves the ability to **construct counterexamples to an argument**. A *counterexample* is a possible situation in which the premises are true and the conclusion isn’t. Any argument that does not admit any counterexamples is called a *valid* argument.

*VALID*:

The U.S. should import Chinese products only if China regulates those products.

China does not strictly regulate its products.

∴ The U.S. should not import those products.

*Invalid* arguments thus have possible situations in which the premises are true and the conclusion isn’t. What we mean by “possible” here should be understood as meaning “not yielding a contradiction.” Thus, something that is both a circle and not a circle is a logical impossibility, but Khalifa spontaneously turning into a jelly donut is a logical possibility.

However, in certain contexts, it’s often helpful to not merely assert that there is no contradiction when the premises are true and the conclusion is false. People often want a

rationale—an explanation as to how it is possible that the premises are true and the conclusion is false. For this reason, **when constructing a counterexample, you should:**

1. **Affirm of all the argument’s premises.**
2. **Deny the argument’s conclusion.**
3. **An explanation of *how this is possible*—that is, how the conclusion can still be untrue while the premises are all true.**

For example, take the following invalid argument.

If China strictly regulates its products, then the U.S. should import those products.  
 China does not strictly regulate its products.  
 ∴ The U.S. should not import those products.

Let’s go through our three rules:

Suppose that goods produced in China are inspected and regulated in the U.S. (Answers the *How possible?* question). In this case, China does not regulate its products (Affirms Premise 2). However, it would nevertheless be true that if the Chinese were to perform their own regulations, the U.S. should import their products (Affirms Premise 1). Nevertheless, this would provide a good reason to import their products, on the assumption that production costs remain low (Denies Conclusion).

Using this example, let’s consider criteria that make for a good explanation/answer to the *how possible?* question (criterion #3 in our “Counterexample” checklist). **In constructing an explanation for a counterexample, keep the following in mind:**

- 1) **The explanation need not be true; it only needs to be conceivable.** While there has been talk of U.S. companies being held accountable for safety-checks of Chinese parts and products that they import, it is not actually the case.
- 2) **However, an explanation that is closer to reality usually illustrates the force of a counterexample more vividly than one that requires extravagant leaps in imagination.** For example, we could have used the following:
  - a. Suppose that goods produced in China are inspected and regulated by little green men inhabiting the moon...(etc.)

*Logically* speaking, we can construct a counterexample just as easily using this scenario, but a typical person will start focusing on how improbable it is that little green men inhabit the moon rather than on the invalidity of their argument. Of course, sometimes you can’t avoid wacky counterexamples...usually when the premises are very wacky.

- 3) **Stay as close to the premises and conclusion as you can in constructing a counterexample.** The more information you need to add to your “story” or explanation as to how the premises can be true and the conclusion false, the more likely your audience is to get bogged down in the details and forget that the issue is the invalidity of an argument. For example, the following counterexample taxes a typical person’s ability and patience to appreciate the invalidity of the argument we’re evaluating:
  - a. Suppose that Chiang Kai-shek, rather than Mao Zedong, seized control of the Chinese Mainland, and that as a result of renouncing Communism, China had more extensive and collegial relations with the West during the Cold War.

As a result, they formed a partnership with a French regulation company named *Régulateur*...

Logically speaking, this works just as well as our example, but just like with the little green man example, you've bogged down your audience with details that distract from the main point: that there's a bad argument being offered.

**You can write a counterexample in the following format:**

***Suppose that [INSERT ANSWER TO HOW POSSIBLE QUESTION HERE]. Then it could still be the case that [AFFIRM PREMISES HERE], but nevertheless [DENY CONCLUSION HERE].***

For example:

*Suppose that the most affordable goods are produced in China and are inspected and regulated in the U.S. Then it could still be the case that China does not regulate its products and that if the Chinese were to perform their own regulations, the U.S. should import their products, but the U.S. should nevertheless import Chinese products.*

There is a very important point to underscore here. Note that we had two arguments, one valid and one invalid, with exactly the same conclusion, namely that the U.S. should not import Chinese products. What this shows is that **you can agree with an author's conclusion, while at the same time thinking the author's reasoning is flawed.** What does this mean for you? First, agreement of opinion isn't enough. Suppose that two people, Alice and Ben, vote for the same candidate. Alice votes for the candidate because his voting record is consistent with her political values, which include gender and racial equity; Ben is a bigot, who votes for the candidate because he is the only white man running for office. Clearly, the two of them reach the same conclusion—that the candidate in question is the best—for very different reasons. Obviously, the difference can be more subtle, and this where you have to be more nuanced in your analysis. But the chief point is that **you should accept conclusions for the best reasons available, not simply because you've found someone you agree with.**

### ***Induction, deduction, and counterexamples***

In our table of common argument forms, we distinguished *inductive* from *deductive* arguments. Deductive arguments are valid; they admit no counterexamples. However, there are many arguments that admit counterexamples that we nevertheless consider to be quite good. For example, the following is a generally regarded as a good inference though it is not deductively valid:

My parents have told me that my name is Kareem Khalifa.

∴ My name is Kareem Khalifa.

However, the conclusion can be false and the premises true. For example, suppose my parents are deceitful people and really named me Bobo Laughingstock, forging documents at every possible turn to hide my true name from me. This is an *inductive* argument (an Argument from Authority in the table above). All inductive arguments admit counterexamples.

Now, clearly we would be fairly annoyed if people raised Bobo-Laughingstock types of counterexamples to such banal inferences as the one concerning our names. This is why I like to distinguish between *counterexamples* and *defeaters* (see the table). Defeaters are the kinds of counterexamples we should take seriously. A defeater is a counterexample that takes the aforementioned constraints on explanations very, very seriously, and is generally a plausible challenge to an inductive argument. The two rightmost columns on the table, labeled *Defeaters* and *Rules*, are designed to help you construct defeaters. Wherever possible, you should **investigate whether an argument succumbs to a defeater**, and you should also **ascertain if the author has not observed one of the rules**. These are very effective ways of seeing if the premises support the conclusion.

## Are the premises true?

However, sometimes we can have valid arguments that still aren't very good, e.g.,

If we had infinite resources, we could end world hunger.

We have infinite resources.

∴ We can end world hunger.

This is just a modus ponens, the paradigm case of a valid argument. Thus, it is not possible that the premises are true when the conclusion is false. However, it's still possible that the premises are false. (Compare: it is not possible for a car without gasoline to move, but it is possible for a car to be without gasoline). This is clearly the case with the second premise: we do not have infinite resources.

Of course, most interesting pieces of nonfiction rest on premises that are more difficult to criticize than this. For example, consider the following argument:

If murder is wrong, then eating meat is wrong.

Murder is wrong.

∴ Eating meat is wrong.

This is a valid argument, but I suspect that many of think that this is not a very good argument. (Even vegetarians may accept the conclusion, while denying the premises. As before, you should accept a conclusion for the best reasons available, not simply because you agree with the conclusion.) Here are some potential diagnoses. The first premise *may* very well be false. How do we discern this? We **pay close attention to a premise's logical vocabulary**, in this case "if ... then." By paying close attention to this vocabulary, we can **ascertain the conditions under which a premise will be false**. Here is a list, inspired partly by the preceding argument forms:

Statement form	Conditions when false
If $p$ , then $q$ .	$p$ is true and $q$ is false
Either $p$ or $q$ .	Both $p$ and $q$ are false.
Both $p$ and $q$ .	Either $p$ is false or $q$ is false.
$p$ if and only if $q$ .	Either $p$ is true and $q$ is false; or $p$ is false and $q$ is true.
All $F$ 's are $G$ 's.	At least one $F$ is not a $G$ .
No $F$ 's are $G$ 's.	At least one $F$ is a $G$ .
Some $F$ 's is a $G$ .	No $F$ 's are $G$ 's.

In this case, the statement will be false when murder is wrong and eating meat is right. But wouldn't more people agree to this than to the vegetarian's claim that both murder and eating meat are wrong? While it certainly isn't the case that majority opinion is always right, it is fair to say that you should **look out for controversial assumptions**. What counts as controversial will depend on the author's audience, and in academic settings, this is often specialists in his/her field. For example, there is a philosophical consensus that Descartes was a rationalist, so a person is free to disagree with that consensus, but only if he/she provides an *argument* to that effect. If she merely takes it up as a brute assumption that Descartes was not a rationalist, she can be rightly criticized. So more generally, **controversial arguments are interesting; controversial assumptions are criticizable**.

If you find a controversial assumption, a strategic divide looms: you can **either push the burden of proof on your opponent or seek to defend the default opinion**.

Returning to our meat-eating example, the first horn of this strategy might involve asking your opponent to provide reasons for why murder and eating meat are similar. **Shifting the burden of proof should take the form of a question, particularly in verbal exchanges**, e.g., "Why do you think that the wrongness of meat-eating follows from the wrongness of murder? This isn't immediately obvious to me." While there are certainly some issues of etiquette here, I think that there's also an important intellectual virtue at play in asking a question rather than just telling people that you think their controversial assumption is wrong. Namely, asking questions is a sign of an open mind, and a much more effective way to discover the best reasons for a position than simply naysaying what you find initially implausible. To be sure, these are *pointed* questions, but that's also part of being a *critical* thinker<sup>1</sup>.

The second horn of the strategy would involve providing an argument why murder and eating meat are different. Here you should **construct arguments using the basic argument patterns** above, and **these arguments should have the denial of the author's controversial premise as their conclusion**. I call these *counterarguments*. For example, we might reply in the following manner:

If an act harms a rational being, then it is wrong.  
Murder is an act that harms a rational being.  
∴ Murder is wrong.

If an act is wrong, then it harms a rational being.  
If eating animals is wrong, then it harms animals.  
Animals are not rational beings.  
∴ So eating animals is not wrong.

In this way, we offer a principled counterargument as to the author's initial premise "If murder is wrong, then eating meat is wrong." Note that this argument itself might be criticized. For example, murdering children, the insane, and mentally challenged people is clearly wrong, though these are not entirely rational people. So the second premise of the first counterargument is controversial. Can you find any other premises that are debatable?

We could imagine our vegetarian friend challenging nearly all of the premises in these two arguments. Perhaps she would then offer her own counterarguments, and these would rest on controversial premises for which we would then provide counterarguments, and so on, *ad nauseum*. There just may not be sufficient overlap of uncontroversial

---

<sup>1</sup> While I won't make much of it here, note that by raising a question, you've also raised a new *issue*, and thus opened up a new venue to apply your critical thinking skills.

assumptions to settle the matter. Nevertheless, there are two important things to note. First, **even if there is no resolution to a debate, your ideas are sharpened by engaging smart people with very different assumptions than you.** We can imagine that a vegetarian who presses us on these counterarguments will force us to think long and hard about what's wrong, rational, etc.

Second, there are certain stopping points. First, **if an interlocutor contradicts him/herself, his/her position is untenable.** So always look for contradictions. Second, there are standard forms of retrenchment that come up in these kinds of exchanges, so **beware of the following:**

- 1) **Retrenched positions that are so weak that they are compatible with, or entailed by the criticisms leveled against them.** Often, when challenged, people will concede certain points without admitting that they're conceding certain points. For example, suppose that in response to shifting the burden of proof, our vegetarian friend says, "Well both murder and meat-eating involve killing innocents," to which you reply, "Certain medicines kill innocent microbes," to which she replies, "Yes, but microbes don't count." But wasn't your initial question exactly designed to target the claim that non-human animals should "count"? So something is remiss with the vegetarian's responses in this exchange.
- 2) **Retrenched positions that use vague language to avoid being refuted.** Imagine a different response to the burden of proof; the vegetarian says "All life is sacred." You run your medicine and microbe argument by her, but this time she says, "Well microbes aren't the kind of 'life' I'm talking about." At this point, you should push the burden of proof even more strongly on this person, "What do you mean by 'life' then?"
- 3) **Retrenched positions that ignore or dismiss well-established empirical facts.** Well-established empirical facts are typically seen as providing the stopping points for most controversies. It is just undeniable that certain things happen in a laboratory test, or that a certain study reveals a very stable set of results. Short of being an expert in experimental design, you should take these studies seriously—unless you have strong evidence of one of the aforementioned inductive defeaters. For example, looking at our second counterargument, a vegetarian might cite the various studies on primate cognition to deny that animals are rational. However, note that this would only mean that we shouldn't eat *primates*, but chickens and cows might be perfectly alright if they lack these cognitive capacities. This points to another rule of thumb: **be sensitive to how strongly a counterargument refutes a targeted premise**, regardless of whether the counterargument or the premise is what you endorse.

## Summary of rules

Well that just about sums it up. There are a lot of rules and concepts here, and you won't perfect critical, philosophical reading overnight. This leads to the last rule: **keep practicing reading with these rules in mind.** They will eventually become second nature. Here they are in list-like form:

*The most basic rule of all:*

1. When taking notes on a text, answer these six questions:
  - a. What is the issue?
  - b. What is the conclusion or thesis?
  - c. What are the premises?
  - d. How, by the author's lights, do the premises support the conclusion?
  - e. Do the premises adequately support the conclusion?
  - f. Are the premises true?

*What is the issue?*

2. Do a quick skim to identify the issue before doing a closer reading.
3. Good writers will almost always introduce the issue in the first few sentences of their works.
4. Issues are often things of which the author takes pains to underscore their importance.
5. Issues are often things that the author treats as in tension or problematic in some sense.
6. Phrase your issue as a question, not as a word or a phrase.
7. Look out for when authors change the issue.
8. Make your question as precise as possible.
9. Wherever possible, frame the issue as a yes-no question.

*What is the conclusion?*

10. Go back and forth between formulating tentative guesses about the issue and formulating the conclusion, until you have the most precise formulation of both.
11. Always look for conclusion-indicators, i.e., words such as "therefore," "thus," "hence," "so," "as a result," "it follows that," etc.
12. Not all conclusions are indicated.
13. Be careful not to mistake intermediate conclusions for ultimate conclusions.

*What are the premises?*

14. Ascertain an author's reasons before you criticize them.
15. Avoid straw man arguments.
16. Be charitable.
17. Always identify the conclusion before you identify the premises.
18. Look for premise-indicators, e.g., "since," "because," "it follows from," "as supported by," etc.
19. Not all premises will be indicated as such.
20. Distinguish details from what is absolutely essential to supporting the conclusion.
21. To ascertain what is necessary, engage in critical imaginative steps.
22. State premises as generally as possible without being too vague.
23. Don't lose sight of the conclusion when identifying the author's premises.

*How, by the author's lights, do the premises support the conclusion?*

24. Get familiar with different patterns of argument.
25. Paraphrase parts of passages as instances of these forms of argument.
26. Where appropriate, paraphrase arguments such that independent premises are carefully distinguished from dependent premises.

27. Recognize the larger argument of which intermediate conclusions are a part or else you'll lose the forest for the trees.
28. Unless the intermediate conclusion *really, really* helps you to understand the flow of the argument, leave it out.
29. When paraphrasing:
  - a. List the premises in an *order* which makes the structure of the argument clear, minimally in *standard form*;
  - b. *Simplify* the language of the original text by trading out more elliptical and counterintuitive language for more concrete and concise language;
  - c. *Eliminate irrelevant propositions*;
  - d. Provide *uniformity* of terms and language;
  - e. Make *implicit premises* explicit, provided they are necessary to get an argument to fit one of the aforementioned forms.

*Do the premises adequately support the conclusion?*

30. Construct counterexamples to an argument.
31. When constructing a counterexample, you should:
  - a. Affirm of all the argument's premises.
  - b. Deny the argument's conclusion.
  - c. An explanation of *how this is possible*—that is, how the conclusion can still be untrue while the premises are all true.
32. In constructing an explanation for a counterexample, keep the following in mind:
  - a. The explanation need not be true; it only needs to be conceivable.
  - b. However, an explanation that is closer to reality usually illustrates the force of a counterexample more vividly than one that requires extravagant leaps in imagination.
  - c. Stay as close to the premises and conclusion as you can in constructing a counterexample.
33. You can write a counterexample in the following format:  
*Suppose that* [INSERT ANSWER TO HOW POSSIBLE QUESTION HERE]. *Then it could still be the case that* [AFFIRM PREMISES HERE], *but nevertheless* [DENY CONCLUSION HERE].
34. You can agree with an author's conclusion, while at the same time thinking the author's reasoning is flawed.
35. You should accept conclusions for the best reasons available, not simply because you've found someone you agree with.
36. Investigate whether an inductive argument succumbs to a defeater.
37. Ascertain if the author has not observed one of the rules for inductive inference.

*Are the premises true?*

38. Pay close attention to a premise's logical vocabulary.
39. Ascertain the conditions under which a premise will be false.
40. Look out for controversial assumptions.
41. Controversial *arguments* are interesting; controversial *assumptions* are criticizable.
42. When faced with a controversial assumption, either push the burden of proof on your opponent or seek to defend the default opinion.
43. Shifting the burden of proof should take the form of a question, particularly in verbal exchanges.



44. Construct counterarguments to basic assumptions using the basic argument patterns, in which denial of the author's controversial premise is the conclusion.
45. Even if there is no resolution to a debate, your ideas are sharpened by engaging smart people with very different assumptions than you.
46. If an interlocutor contradicts him/herself, his/her position is untenable.
47. Beware of the following:
  - a. Retrenched positions that are so weak that they are compatible with, or entailed by the criticisms leveled against them;
  - b. Retrenched positions that use vague language to avoid being refuted; and
  - c. Retrenched positions that ignore or dismiss well-established empirical facts.
48. Keep practicing reading with these rules in mind.