


# Introduction to Logic and Critical Thinking

Version 1.4



Matthew J. Van Cleave  
Lansing Community College



Introduction to Logic and Critical Thinking by Matthew J. Van Cleave is licensed under a Creative Commons Attribution 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

### **4.1 Formal vs. informal fallacies**

A **fallacy** is simply a mistake in reasoning. Some fallacies are formal and some are informal. In chapter 2, we saw that we could define validity formally and thus could determine whether an argument was valid or invalid without even having to know or understand what the argument was about. We saw that we could define certain valid rules of inference, such as modus ponens and modus tollens. These inference patterns are valid in virtue of their form, not their content. That is, any argument that has the same form as modus ponens or modus tollens will automatically be valid. A formal fallacy is simply an argument whose form is invalid. Thus, any argument that has that form will automatically be invalid, regardless of the meaning of the sentences. Two formal fallacies that are similar to, but should never be confused with, modus ponens and modus tollens are **denying the antecedent** and **affirming the consequent**. Here are the forms of those invalid inferences:

#### Denying the antecedent

$$\begin{array}{l} p \supset q \\ \sim p \\ \therefore \sim q \end{array}$$

#### Affirming the consequent

$$\begin{array}{l} p \supset q \\ q \\ \therefore p \end{array}$$

Any argument that has either of these forms is an invalid argument. For example:

1. If Kant was a deontologist, then he was a non-consequentialist.
2. Kant was not a deontologist.
3. Therefore, Kant was not a non-consequentialist.

The form of this argument is:

1.  $D \supset C$
2.  $\sim D$
3.  $\therefore \sim C$

As you can see, this argument has the form of the fallacy, denying the antecedent. Thus, we know that this argument is invalid even if we don't know what "Kant" or "deontologist" or "non-consequentialist" means. ("Kant" was a famous German philosopher from the early 1800s, whereas "deontology" and "non-consequentialist" are terms that come from ethical theory.) It is mark of a formal fallacy that we can identify it even if we don't really understand the meanings of the sentences in the argument. Recall our Jabberwocky argument from chapter 2. Here's an argument which uses silly, made-up words from Lewis Carroll's "Jabberwocky." See if you can determine whether the argument's form is valid or invalid:

1. If toves are brillig then toves are slithy.
2. Toves are slithy
3. Therefore, toves are brillig.

You should be able to see that this argument has the form of affirming the consequent:

1.  $B \supset S$
2.  $S$
3.  $\therefore B$

As such, we know that the argument is invalid, even though we haven't got a clue what "toves" are or what "slithy" or "brillig" means. The point is that we can identify formal fallacies without having to know what they mean.

In contrast, **informal fallacies** are those which cannot be identified without understanding the concepts involved in the argument. A paradigm example of an informal fallacy is the fallacy of composition. We will consider this fallacy in the next sub-section. In the remaining subsections, we will consider a number of other informal logical fallacies.

#### 4.1.1 Composition fallacy

Consider the following argument:

Each member on the gymnastics team weighs less than 110 lbs.  
Therefore, the whole gymnastics team weighs less than 110 lbs.

This argument commits the composition fallacy. In the composition fallacy one argues that since each part of the whole has a certain feature, it follows that the whole has that same feature. However, you cannot generally identify any argument that moves from statements about parts to statements about wholes as committing the composition fallacy because whether or not there is a fallacy depends on what feature we are attributing to the parts and wholes. Here is an example of an argument that moves from claims about the parts possessing a feature to a claim about the whole possessing that same feature, but doesn't commit the composition fallacy:

Every part of the car is made of plastic. Therefore, the whole car is made of plastic.

This conclusion *does* follow from the premises; there is no fallacy here. The difference between this argument and the preceding argument (about the gymnastics team) isn't their form. In fact both arguments have the same form:

Every part of X has the feature f. Therefore, the whole X has the feature f.

And yet one of the arguments is clearly fallacious, while the other isn't. The difference between the two arguments is not their form, but their content. That is, the difference is what feature is being attributed to the parts and wholes. Some features (like *weighing a certain amount*) are such that if they belong to each part, then it does not follow that they belong to the whole. Other features (such as *being made of plastic*) are such that if they belong to each part, it follows that they belong to the whole.

Here is another example:

Every member of the team has been to Paris. Therefore the team has been to Paris.

The conclusion of this argument does not follow. Just because each member of the team has been to Paris, it doesn't follow that the whole team has been to Paris, since it may not have been the case that each individual was there at the same time and was there in their capacity as a member of the team. Thus, even though it is plausible to say that the team is composed of every member of the team, it doesn't follow that since every member of the team has been to Paris, the whole team has been to Paris. Contrast that example with this one:

Every member of the team was on the plane. Therefore, the whole team was on the plane.

This argument, in contrast to the last one, contains no fallacy. It is true that if every member is on the plane then the whole team is on the plane. And yet these two arguments have almost exactly the same form. The only difference is that the first argument is talking about the property, *having been to Paris*, whereas the second argument is talking about the property, *being on the plane*. The only reason we are able to identify the first argument as committing the composition fallacy and the second argument as not committing a fallacy is that we understand the relationship between the concepts involved. In the first case, we understand that it is possible that every member could have been to Paris without the team ever having been; in the second case we understand that as long as every member of the team is on the plane, it has to be true that the whole team is on the plane. The take home point here is that in order to identify whether an argument has committed the composition fallacy, one must understand the concepts involved in the argument. This is the mark of an informal fallacy: we have to rely on our understanding of the meanings of the words or concepts involved, rather than simply being able to identify the fallacy from its form.

#### **4.1.2 Division fallacy**

The division fallacy is like the composition fallacy and they are easy to confuse. The difference is that the division fallacy argues that since the whole has some feature, each part must also have that feature. The composition fallacy, as we have just seen, goes in the opposite direction: since each part has some feature, the whole must have that same feature. Here is an example of a division fallacy:

The house costs 1 million dollars. Therefore, each part of the house costs 1 million dollars.

This is clearly a fallacy. Just because the whole house costs 1 million dollars, it doesn't follow that each part of the house costs 1 million dollars. However, here is an argument that has the same form, but that doesn't commit the division fallacy:

The whole team died in the plane crash. Therefore each individual on the team died in the plane crash.

In this example, since we seem to be referring to one plane crash in which all the members of the team died ("the" plane crash), it follows that if the whole team died in the crash, then every individual on the team died in the crash. So this argument does not commit the division fallacy. In contrast, the following argument has exactly the same form, but does commit the division fallacy:

The team played its worst game ever tonight. Therefore, each individual on the team played their worst game ever tonight.

It can be true that the whole team played its worst game ever even if it is true that no individual on the team played their worst game ever. Thus, this argument does commit the fallacy of division even though it has the same form as the previous argument, which doesn't commit the fallacy of division. This shows (again) that in order to identify informal fallacies (like composition and division), we must rely on our understanding of the concepts involved in the argument. Some concepts (like "team" and "dying in a plane crash") are such that if they apply to the whole, they also apply to all the parts. Other concepts (like "team" and "worst game played") are such that they can apply to the whole even if they do not apply to all the parts.

### **4.1.3 Begging the question**

Consider the following argument:

Capital punishment is justified for crimes such as rape and murder because it is quite legitimate and appropriate for the state to put to death someone who has committed such heinous and inhuman acts.

The premise indicator, "because" denotes the premise and (derivatively) the conclusion of this argument. In standard form, the argument is this:

1. It is legitimate and appropriate for the state to put to death someone who commits rape or murder.
2. Therefore, capital punishment is justified for crimes such as rape and murder.

You should notice something peculiar about this argument: the premise is essentially the same claim as the conclusion. The only difference is that the premise spells out what capital punishment means (the state putting criminals to death) whereas the conclusion just refers to capital punishment by name, and the premise uses terms like “legitimate” and “appropriate” whereas the conclusion uses the related term, “justified.” But these differences don’t add up to any real differences in meaning. Thus, the premise is essentially saying the same thing as the conclusion. This is a problem: we want our premise to provide a reason for accepting the conclusion. But if the premise is the same claim as the conclusion, then it can’t possibly provide a reason for accepting the conclusion! **Begging the question** occurs when one (either explicitly or implicitly) assumes the truth of the conclusion in one or more of the premises. Begging the question is thus a kind of circular reasoning.

One interesting feature of this fallacy is that *formally* there is nothing wrong with arguments of this form. Here is what I mean. Consider an argument that explicitly commits the fallacy of begging the question. For example,

1. Capital punishment is morally permissible
2. Therefore, capital punishment is morally permissible

Now, apply any method of assessing validity to this argument and you will see that it is valid by any method. If we use the informal test (by trying to imagine that the premises are true while the conclusion is false), then the argument passes the test, since any time the premise is true, the conclusion will have to be true as well (since it is the exact same statement). Likewise, the argument is valid by our formal test of validity, truth tables. But while this argument is technically valid, it is still a really bad argument. Why? Because the point of giving an argument in the first place is to provide some reason for thinking the conclusion is true *for those who don’t already accept the conclusion*. But if one doesn’t already accept the conclusion, then simply restating the conclusion in a different way isn’t going to convince them. Rather, a good argument will provide some reason for accepting the conclusion that is sufficiently independent of that conclusion itself. Begging the question utterly fails to do this and this is why it counts as an informal fallacy. What is interesting about begging the question is that there is absolutely nothing wrong with the argument formally.



Whether or not an argument begs the question is not always an easy matter to sort out. As with all informal fallacies, detecting it requires a careful understanding of the meaning of the statements involved in the argument. Here is an example of an argument where it is not as clear whether there is a fallacy of begging the question:

Christian belief is warranted because according to Christianity there exists a being called “the Holy Spirit” which reliably guides Christians towards the truth regarding the central claims of Christianity.<sup>1</sup>

One might think that there is a kind of circularity (or begging the question) involved in this argument since the argument appears to assume the truth of Christianity in justifying the claim that Christianity is true. But whether or not this argument really does beg the question is something on which there is much debate within the sub-field of philosophy called epistemology (“study of knowledge”). The philosopher Alvin Plantinga argues persuasively that the argument does not beg the question, but being able to assess that argument takes patient years of study in the field of epistemology (not to mention a careful engagement with Plantinga’s work). As this example illustrates, the issue of whether an argument begs the question requires us to draw on our general knowledge of the world. This is the mark of an informal, rather than formal, fallacy.

#### **4.1.4 False dichotomy**

Suppose I were to argue as follows:

Raising taxes on the wealthy will either hurt the economy or it will help it.  
But it won’t help the economy. Therefore it will hurt the economy.

The standard form of this argument is:

1. Either raising taxes on the wealthy will hurt the economy or it will help it.
2. Raising taxes on the wealthy won’t help the economy.
3. Therefore, raising taxes on the wealthy will hurt the economy.

---

<sup>1</sup> This is a much simplified version of the view defended by Christian philosophers such as Alvin Plantinga. Plantinga defends (something like) this claim in: Plantinga, A. 2000. *Warranted Christian Belief*. Oxford, UK: Oxford University Press.

This argument contains a fallacy called a “false dichotomy.” A **false dichotomy** is simply a disjunction that does not exhaust all of the possible options. In this case, the problematic disjunction is the first premise: either raising the taxes on the wealthy will hurt the economy or it will help it. But these aren’t the only options. Another option is that raising taxes on the wealthy will have no effect on the economy. Notice that the argument above has the form of a disjunctive syllogism:

$$\begin{array}{l} A \vee B \\ \sim A \\ \therefore B \end{array}$$

However, since the first premise presents two options as if they were the only two options, when in fact they aren’t, the first premise is false and the argument fails. Notice that the form of the argument is perfectly good—the argument is valid. The problem is that this argument isn’t sound because the first premise of the argument commits the false dichotomy fallacy. False dichotomies are commonly encountered in the context of a disjunctive syllogism or constructive dilemma (see chapter 2).

In a speech made on April 5, 2004, President Bush made the following remarks about the causes of the Iraq war:

Saddam Hussein once again defied the demands of the world. And so I had a choice: Do I take the word of a madman, do I trust a person who had used weapons of mass destruction on his own people, plus people in the neighborhood, or do I take the steps necessary to defend the country? Given that choice, I will defend America every time.

The false dichotomy here is the claim that:

Either I trust the word of a madman or I defend America (by going to war against Saddam Hussein’s regime).

The problem is that these aren’t the only options. Other options include ongoing diplomacy and economic sanctions. Thus, even if it true that Bush shouldn’t have trusted the word of Hussein, it doesn’t follow that the only other

option is going to war against Hussein's regime. (Furthermore, it isn't clear in what sense this was needed to defend *America*.) That is a false dichotomy.

As with all the previous informal fallacies we've considered, the false dichotomy fallacy requires an understanding of the concepts involved. Thus, we have to use our understanding of world in order to assess whether a false dichotomy fallacy is being committed or not.

#### **4.1.5 Equivocation**

Consider the following argument:

Children are a headache. Aspirin will make headaches go away.  
Therefore, aspirin will make children go away.

This is a silly argument, but it illustrates the fallacy of equivocation. The problem is that the word "headache" is used equivocally—that is, in two different senses. In the first premise, "headache" is used figuratively, whereas in the second premise "headache" is used literally. The argument is only successful if the meaning of "headache" is the same in both premises. But it isn't and this is what makes this argument an instance of the fallacy of equivocation.

Here's another example:

Taking a logic class helps you learn how to argue. But there is already too much hostility in the world today, and the fewer arguments the better.  
Therefore, you shouldn't take a logic class.

In this example, the word "argue" and "argument" are used equivocally. Hopefully, at this point in the text, you recognize the difference. (If not, go back and reread section 1.1.)

The fallacy of equivocation is not always so easy to spot. Here is a trickier example:

The existence of laws depends on the existence of intelligent beings like humans who create the laws. However, some laws existed before there were any humans (e.g., laws of physics). Therefore, there must be some non-human, intelligent being that created these laws of nature.

The term “law” is used equivocally here. In the first premise it is used to refer to societal laws, such as criminal law; in the second premise it is used to refer to laws of nature. Although we use the term “law” to apply to both cases, they are importantly different. Societal laws, such as the criminal law of a society, are enforced by people and there are punishments for breaking the laws. Natural laws, such as laws of physics, cannot be broken and thus there are no punishments for breaking them. (Does it make sense to scold the electron for not doing what the law says it will do?)

As with every informal fallacy we have examined in this section, equivocation can only be identified by understanding the meanings of the words involved. In fact, the definition of the fallacy of equivocation refers to this very fact: the same word is being used in two different senses (i.e., with two different meanings). So, unlike formal fallacies, identifying the fallacy of equivocation requires that we draw on our understanding of the meaning of words and of our understanding of the world, generally.

## **4.2 Slippery slope fallacies**

Slippery slope fallacies depend on the concept of **vagueness**. When a concept or claim is vague, it means that we don’t know precisely what claim is being made, or what the boundaries of the concept are. The classic example used to illustrate vagueness is the “**sorites paradox**.” The term “sorites” is the Greek term for “heap” and the paradox comes from ancient Greek philosophy. Here is the paradox. I will give you two claims that each sound very plausible, but in fact lead to a paradox. Here are the two claims:

1. One grain of sand is not a heap of sand.
2. If I start with something that is not a heap of sand, then adding one grain of sand to that will not create a heap of sand.

For example, two grains of sand is not a heap, thus (by the second claim) neither is three grains of sand. But since three grains of sand is not a heap then (by the second claim again) neither is four grains of sand. You can probably see where this is going. By continuing to add one grain of sand over and over, I will eventually end up with something that is clearly a heap of sand, but that won’t be counted as a heap of sand if we accept both claims 1 and 2 above.

Philosophers continue to argue and debate about how to resolve the sorites paradox, but the point for us is just to illustrate the concept of vagueness. The concept “heap” is a vague concept in this example. But so are so many other concepts, such as color concepts (red, yellow, green, etc.), moral concepts (right, wrong, good, bad), and just about any other concept you can think of. The one domain that seems to be unaffected by vagueness is mathematical and logical concepts. There are two fallacies related to vagueness: the causal slippery slope and the conceptual slippery slope. We’ll cover the conceptual slippery slope first since it relates most closely to the concept of vagueness I’ve explained above.

### **4.2.1 Conceptual slippery slope**

It may be true that there is no essential difference between 499 grains of sand and 500 grains of sand. But even if that is so, it doesn’t follow that there is no difference between 1 grain of sand and 5 billion grains of sand. In general, just because we cannot draw a distinction between A and B, and we cannot draw a distinction between B and C, it doesn’t mean we cannot draw a distinction between A and C. Here is an example of a conceptual slippery slope fallacy.

It is illegal for anyone under 21 to drink alcohol. But there is no difference between someone who is 21 and someone who is 20 years 11 months old. So there is nothing wrong with someone who is 20 years and 11 months old drinking. But since there is no real distinction between being one month older and one month younger, there shouldn’t be anything wrong with drinking at any age. Therefore, there is nothing wrong with allowing a 10 year old to drink alcohol.

Imagine the life of an individual in stages of 1 month intervals. Even if it is true that there is no distinction in kind between any one of those stages, it doesn’t follow that there isn’t a distinction to be drawn at the extremes of either end. Clearly there is a difference between a 5 year old and a 25 year old—a distinction in kind that is relevant to whether they should be allowed to drink alcohol. The conceptual slippery slope fallacy assumes that because we cannot draw a distinction between adjacent stages, we cannot draw a distinction at all between any stages. One clear way of illustrating this is with color. Think of a color spectrum from purple to red to orange to yellow to green to blue. Each color grades into the next without there being any distinguishable boundaries

between the colors—a continuous spectrum. Even if it is true that for any two adjacent hues on the color wheel, we cannot distinguish between the two, it doesn't follow from this that there is no distinction to be drawn between any two portions of the color wheel, because then we'd be committed to saying that there is no distinguishable difference between purple and yellow! The example of the color spectrum illustrates the general point that just because the boundaries between very similar things on a spectrum are vague, it doesn't follow that there are no differences between any two things on that spectrum.

Whether or not one will identify an argument as committing a conceptual slippery slope fallacy, depends on the other things one believes about the world. Thus, whether or not a conceptual slippery slope fallacy has been committed will often be a matter of some debate. It will itself be vague. Here is a good example that illustrates this point.

People are found not guilty by reason of insanity when they cannot avoid breaking the law. But people who are brought up in certain deprived social circumstances are not much more able than the legally insane to avoid breaking the law. So we should not find such individuals guilty any more than those who are legally insane.

Whether there is conceptual slippery slope fallacy here depends on what you think about a host of other things, including individual responsibility, free will, the psychological and social effects of deprived social circumstances such as poverty, lack of opportunity, abuse, etc. Some people may think that there are big differences between those who are legally insane and those who grow up in deprived social circumstances. Others may not think the differences are so great. The issues here are subtle, sensitive, and complex, which is why it is difficult to determine whether there is any fallacy here or not. If the differences between those who are insane and those who are the product of deprived social circumstances turn out to be like the differences between one shade of yellow and an adjacent shade of yellow, then there is no fallacy here. But if the differences turn out to be analogous to those between yellow and green (i.e., with many distinguishable stages of difference between) then there would indeed be a conceptual slippery slope fallacy here. The difficulty of distinguishing instances of the conceptual slippery slope fallacy, and the fact that distinguishing it requires us to draw on our knowledge about the world, shows that the conceptual slippery slope fallacy is an informal fallacy.

### 4.2.2 Causal slippery slope fallacy

The causal slippery slope fallacy is committed when one event is said to lead to some other (usually disastrous) event via a chain of intermediary events. If you have ever seen Direct TV's "get rid of cable" commercials, you will know exactly what I'm talking about. (If you don't know what I'm talking about you should Google it right now and find out. They're quite funny.) Here is an example of a causal slippery slope fallacy (it is adapted from one of the Direct TV commercials):

If you use cable, your cable will probably go on the fritz. If your cable is on the fritz, you will probably get frustrated. When you get frustrated you will probably hit the table. When you hit the table, your young daughter will probably imitate you. When your daughter imitates you, she will probably get thrown out of school. When she gets thrown out of school, she will probably meet undesirables. When she meets undesirables, she will probably marry undesirables. When she marries undesirables, you will probably have a grandson with a dog collar. Therefore, if you use cable, you will probably have a grandson with dog collar.

This example is silly and absurd, yes. But it illustrates the causal slippery slope fallacy. Slippery slope fallacies are always made up of a series of conjunctions of probabilistic conditional statements that link the first event to the last event. A causal slippery slope fallacy is committed when one assumes that just because each individual conditional statement is probable, the conditional that links the first event to the last event is also probable. Even if we grant that each "link" in the chain is individually probable, it doesn't follow that the whole chain (or the conditional that links the first event to the last event) is probable. Suppose, for the sake of the argument, we assign probabilities to each "link" or conditional statement, like this. (I have italicized the consequents of the conditionals and assigned high conditional probabilities to them. The high probability is for the sake of the argument; I don't actually think these things are as probable as I've assumed here.)

If you use cable, then *your cable will probably go on the fritz* (.9)  
 If your cable is on the fritz, then *you will probably get angry* (.9)  
 If you get angry, then *you will probably hit the table* (.9)  
 If you hit the table, *your daughter will probably imitate you* (.8)

If your daughter imitates you, *she will probably be kicked out of school* (.8)

If she is kicked out of school, *she will probably meet undesirables* (.9)

If she meets undesirables, *she will probably marry undesirables* (.8)

If she marries undesirables, *you will probably have a grandson with a dog collar* (.8)

However, even if we grant the probabilities of each link in the chain is high (80-90% probable), the conclusion doesn't even reach a probability higher than chance. Recall that in order to figure the probability of a conjunction, we must multiply the probability of each conjunct:

$$(.9) \times (.9) \times (.9) \times (.8) \times (.8) \times (.9) \times (.8) \times (.8) = .27$$

That means the probability of the conclusion (i.e., that if you use cable, you will have a grandson with a dog collar) is only 27%, despite the fact that each conditional has a relatively high probability! The causal slippery slope fallacy is actually a formal probabilistic fallacy and so could have been discussed in chapter 3 with the other formal probabilistic fallacies. What makes it a formal rather than informal fallacy is that we can identify it without even having to know what the sentences of the argument mean. I could just have easily written out a nonsense argument comprised of series of probabilistic conditional statements. But I would still have been able to identify the causal slippery slope fallacy because I would have seen that there was a series of probabilistic conditional statements leading to a claim that the conclusion of the series was also probable. That is enough to tell me that there is a causal slippery slope fallacy, even if I don't really understand the meanings of the conditional statements.

It is helpful to contrast the causal slippery slope fallacy with the valid form of inference, hypothetical syllogism. Recall that a hypothetical syllogism has the following kind of form:

$A \supset B$

$B \supset C$

$C \supset D$

$D \supset E$

$\therefore A \supset E$



The only difference between this and the causal slippery slope fallacy is that whereas in the hypothetical syllogism, the link between each component is *certain*, in a causal slippery slope fallacy, the link between each event is *probabilistic*. It is the fact that each link is probabilistic that accounts for the fallacy. One way of putting this point is that *probability is not transitive*. Just because A makes B probable and B makes C probable and C makes X probable, it doesn't follow that A makes X probable. In contrast, when the links are certain rather than probable, then if A always leads to B and B always leads to C and C always leads to X, then it has to be the case that A always leads to X.

### **4.3 Fallacies of relevance**

What all fallacies of relevance have in common is that they make an argument or response to an argument that is irrelevant. Fallacies of relevance can be compelling psychologically, but it is important to distinguish between rhetorical techniques that are psychologically compelling, on the one hand, and rationally compelling arguments, on the other. What makes something a fallacy is that it fails to be rationally compelling, once we have carefully considered it. That said, arguments that fail to be rationally compelling may still be psychologically or emotionally compelling. The first fallacy of relevance that we will consider, the ad hominem fallacy, is an excellent example a fallacy that can be psychologically compelling.

#### **4.3.1 Ad hominem**

"Ad hominem" is a Latin phrase that can be translated into English as the phrase, "against the man." In an ad hominem fallacy, instead of responding to (or attacking) the *argument* a person has made, one attacks the *person* him or herself. In short, one attacks the person making the argument rather than the argument itself. Here is an anecdote that reveals an ad hominem fallacy (and that has actually occurred in my ethics class before).

A philosopher named Peter Singer had made an argument that it is morally wrong to spend money on luxuries for oneself rather than give all of your money that you don't strictly need away to charity. The argument is actually an argument from analogy (whose details I discussed in section 3.3), but the essence of the argument is that there are every day in this world children who die preventable deaths, and there are charities who

could save the lives of these children if they are funded by individuals from wealthy countries like our own. Since there are things that we all regularly buy that we don't need (e.g., Starbucks lattes, beer, movie tickets, or extra clothes or shoes we don't really need), if we continue to purchase those things rather than using that money to save the lives of children, then we are essentially contributing to the deaths of those children if we choose to continue to live our lifestyle of buying things we don't need, rather than donating the money to a charity that will save lives of children in need. In response to Singer's argument, one student in the class asked: "Does Peter Singer give his money to charity? Does he do what he says we are all morally required to do?"

The implication of this student's question (which I confirmed by following up with her) was that if Peter Singer himself doesn't donate all his extra money to charities, then his argument isn't any good and can be dismissed. But that would be to commit an *ad hominem* fallacy. Instead of responding to the *argument* that Singer had made, this student attacked *Singer himself*. That is, they wanted to know how Singer lived and whether he was a hypocrite or not. Was he the kind of person who would tell us all that we had to live a certain way but fail to live that way himself? But all of this is irrelevant to assessing Singer's *argument*. Suppose that Singer didn't donate his excess money to charity and instead spent it on luxurious things for himself. Still, the argument that Singer has given can be assessed on its own merits. Even if it were true that Peter Singer was a total hypocrite, his argument may nevertheless be rationally compelling. And it is the quality of the argument that we are interested in, not Peter Singer's personal life and whether or not he is hypocritical. Whether Singer is or isn't a hypocrite, is irrelevant to whether the argument he has put forward is strong or weak, valid or invalid. The argument stands on its own and it is that argument rather than Peter Singer himself that we need to assess.

Nonetheless, there is something psychologically compelling about the question: Does Peter Singer practice what he preaches? I think what makes this question seem compelling is that humans are very interested in finding "cheaters" or hypocrites—those who say one thing and then do another. Evolutionarily, our concern with cheaters makes sense because cheaters can't be trusted and it is essential for us (as a group) to be able to pick out those who can't be trusted. That said, whether or not a person giving an argument is a hypocrite is irrelevant to whether that person's argument is good or bad. So there may be psychological reasons why humans are prone to find certain kinds of ad

hominem fallacies *psychologically compelling*, even though ad hominem fallacies are not *rationally compelling*.

Not every instance in which someone attacks a person's character is an ad hominem fallacy. Suppose a witness is on the stand testifying against a defendant in a court of law. When the witness is cross examined by the defense lawyer, the defense lawyer tries to go for the witness's credibility, perhaps by digging up things about the witness's past. For example, the defense lawyer may find out that the witness cheated on her taxes five years ago or that the witness failed to pay her parking tickets. The reason this isn't an ad hominem fallacy is that in this case the lawyer is trying to establish whether what the witness is saying is true or false and in order to determine that we have to know whether the witness is trustworthy. These facts about the witness's past may be relevant to determining whether we can trust the witness's word. In this case, the witness is making claims that are either true or false rather than giving an argument. In contrast, when we are assessing someone's argument, the argument stands on its own in a way the witness's testimony doesn't. In assessing an argument, we want to know whether the argument is strong or weak and we can evaluate the argument using the logical techniques surveyed in this text. In contrast, when a witness is giving testimony, they aren't trying to argue anything. Rather, they are simply making a claim about what did or didn't happen. So although it may seem that a lawyer is committing an ad hominem fallacy in bringing up things about the witness's past, these things are actually relevant to establishing the witness's credibility. In contrast, when considering an argument that has been given, we don't have to establish the arguer's credibility because we can assess the argument they have given on its own merits. The arguer's personal life is irrelevant.

#### **4.3.2 Straw man**

Suppose that my opponent has argued for a position, call it position A, and in response to his argument, I give a rationally compelling argument against position B, which is related to position A, but is much less plausible (and thus much easier to refute). What I have just done is attacked a straw man—a position that “looks like” the target position, but is actually not that position. When one attacks a straw man, one commits the straw man fallacy. The straw man fallacy misrepresents one's opponent's argument and is thus a kind of irrelevance. Here is an example.

Two candidates for political office in Colorado, Tom and Fred, are having an exchange in a debate in which Tom has laid out his plan for putting more money into health care and education and Fred has laid out his plan which includes earmarking more state money for building more prisons which will create more jobs and, thus, strengthen Colorado's economy. Fred responds to Tom's argument that we need to increase funding to health care and education as follows: "I am surprised, Tom, that you are willing to put our state's economic future at risk by sinking money into these programs that do not help to create jobs. You see, folks, Tom's plan will risk sending our economy into a tailspin, risking harm to thousands of Coloradans. On the other hand, my plan supports a healthy and strong Colorado and would never bet our state's economic security on idealistic notions that simply don't work when the rubber meets the road."

Fred has committed the straw man fallacy. Just because Tom wants to increase funding to health care and education does not mean he does not want to help the economy. Furthermore, increasing funding to health care and education does not entail that fewer jobs will be created. Fred has attacked a position that is not the position that Tom holds, but is in fact a much less plausible, easier to refute position. However, it would be silly for any political candidate to run on a platform that included "harming the economy." Presumably no political candidate would run on such a platform. Nonetheless, this exact kind of straw man is ubiquitous in political discourse in our country.

Here is another example.

Nancy has just argued that we should provide middle schoolers with sex education classes, including how to use contraceptives so that they can practice safe sex should they end up in the situation where they are having sex. Fran responds: "proponents of sex education try to encourage our children to a sex-with-no-strings-attached mentality, which is harmful to our children and to our society."

Fran has committed the straw man (or straw woman) fallacy by misrepresenting Nancy's position. Nancy's position is not that we should encourage children to have sex, but that we should make sure that they are fully informed about sex so that if they do have sex, they go into it at least a little less blindly and are able to make better decision regarding sex.

As with other fallacies of relevance, straw man fallacies can be compelling on some level, even though they are irrelevant. It may be that part of the reason we are taken in by straw man fallacies is that humans are prone to “demonize” the “other”—including those who hold a moral or political position different from our own. It is easy to think bad things about those with whom we do not regularly interact. And it is easy to forget that people who are different than us are still people just like us in all the important respects. Many years ago, atheists were commonly thought of as highly immoral people and stories about the horrible things that atheists did in secret circulated widely. People believed that these strange “others” were capable of the most horrible savagery. After all, they may have reasoned, if you don’t believe there is a God holding us accountable, why be moral? The Jewish philosopher, Baruch Spinoza, was an atheist who lived in the Netherlands in the 17<sup>th</sup> century. He was accused of all sorts of things that were commonly believed about atheists. But he was in fact as upstanding and moral as any person you could imagine. The people who knew Spinoza knew better, but how could so many people be so wrong about Spinoza? I suspect that part of the reason is that since at that time there were very few atheists (or at least very few people actually admitted to it), very few people ever knowingly encountered an atheist. Because of this, the stories about atheists could proliferate without being put in check by the facts. I suspect the same kind of phenomenon explains why certain kinds of straw man fallacies proliferate. If you are a conservative and mostly only interact with other conservatives, you might be prone to holding lots of false beliefs about liberals. And so maybe you are less prone to notice straw man fallacies targeted at liberals because the false beliefs you hold about them incline you to see the straw man fallacies as true.

### **4.3.3 Tu quoque**

“Tu quoque” is a Latin phrase that can be translated into English as “you too” or “you, also.” The tu quoque fallacy is a way of avoiding answering a criticism by bringing up a criticism of your opponent rather than answer the criticism. For example, suppose that two political candidates, A and B, are discussing their policies and A brings up a criticism of B’s policy. In response, B brings up her own criticism of A’s policy rather than respond to A’s criticism of her policy. B has here committed the tu quoque fallacy. The fallacy is best understood as a way of avoiding having to answer a tough criticism that one may not have a good answer to. This kind of thing happens all the time in political discourse.

Tu quoque, as I have presented it, is fallacious when the criticism one raises is simply in order to avoid having to answer a difficult objection to one's argument or view. However, there are circumstances in which a tu quoque kind of response is not fallacious. If the criticism that A brings toward B is a criticism that equally applies not only to A's position but to any position, then B is right to point this fact out. For example, suppose that A criticizes B for taking money from special interest groups. In this case, B would be totally right (and there would be no tu quoque fallacy committed) to respond that not only does A take money from special interest groups, but every political candidate running for office does. That is just a fact of life in American politics today. So A really has no criticism at all to B since everyone does what B is doing and it is in many ways unavoidable. Thus, B could (and should) respond with a "you too" rebuttal and in this case that rebuttal is not a tu quoque fallacy.

#### **4.3.4 Genetic fallacy**

The genetic fallacy occurs when one argues (or, more commonly, implies) that the origin of something (e.g., a theory, idea, policy, etc.) is a reason for rejecting (or accepting) it. For example, suppose that Jack is arguing that we should allow physician assisted suicide and Jill responds that that idea first was used in Nazi Germany. Jill has just committed a genetic fallacy because she is implying that because the idea is associated with Nazi Germany, there must be something wrong with the idea itself. What she should have done instead is explain what, exactly, is wrong with the idea rather than simply assuming that there must be something wrong with it since it has a negative origin. The origin of an idea has nothing inherently to do with its truth or plausibility. Suppose that Hitler constructed a mathematical proof in his early adulthood (he didn't, but just suppose). The validity of that mathematical proof stands on its own; the fact that Hitler was a horrible person has nothing to do with whether the proof is good. Likewise with any other idea: ideas must be assessed on their own merits and the origin of an idea is neither a merit nor demerit of the idea.

Although genetic fallacies are most often committed when one associates an idea with a negative origin, it can also go the other way: one can imply that because the idea has a positive origin, the idea must be true or more plausible. For example, suppose that Jill argues that the Golden Rule is a good way to live one's life because the Golden Rule originated with Jesus in the Sermon on the Mount (it didn't, actually, even though Jesus does state a version of the Golden Rule). Jill has committed the genetic fallacy in assuming that the (presumed)

fact that Jesus is the origin of the Golden Rule has anything to do with whether the Golden Rule is a good idea.

I'll end with an example from William James's seminal work, *The Varieties of Religious Experience*. In that book (originally a set of lectures), James considers the idea that if religious experiences could be explained in terms of neurological causes, then the legitimacy of the religious experience is undermined. James, being a materialist who thinks that all mental states are physical states—ultimately a matter of complex brain chemistry, says that the fact that any religious experience has a physical cause does not undermine that veracity of that experience. Although he doesn't use the term explicitly, James claims that the claim that the physical origin of some experience undermines the veracity of that experience is a genetic fallacy. Origin is irrelevant for assessing the veracity of an experience, James thinks. In fact, he thinks that religious dogmatists who take the origin of the Bible to be the word of God are making exactly the same mistake as those who think that a physical explanation of a religious experience would undermine its veracity. We must assess ideas for their merits, James thinks, not their origins.

#### **4.3.5 Appeal to consequences**

The appeal to consequences fallacy is like the reverse of the genetic fallacy: whereas the genetic fallacy consists in the mistake of trying to assess the truth or reasonableness of an idea based on the *origin* of the idea, the appeal to consequences fallacy consists in the mistake of trying to assess the truth or reasonableness of an idea based on the (typically negative) *consequences* of accepting that idea. For example, suppose that the results of a study revealed that there are IQ differences between different races (this is a fictitious example, there is no such study that I know of). In debating the results of this study, one researcher claims that if we were to accept these results, it would lead to increased racism in our society, which is not tolerable. Therefore, these results must not be right since if they were accepted, it would lead to increased racism. The researcher who responded in this way has committed the appeal to consequences fallacy. Again, we must assess the study on its own merits. If there is something wrong with the study, some flaw in its design, for example, then that would be a relevant criticism of the study. However, the fact that the results of the study, if widely circulated, would have a negative effect on society is not a reason for rejecting these results as false. The consequences of some idea (good or bad) are irrelevant to the truth or reasonableness of that idea.

Notice that the researchers, being convinced of the negative consequences of the study on society, might rationally choose *not to publish* the study (for fear of the negative consequences). This is totally fine and is not a fallacy. The fallacy consists not in choosing not to publish something that could have adverse consequences, but in claiming that the *results themselves* are undermined by the negative consequences they could have. The fact is, sometimes truth can have negative consequences and falsehoods can have positive consequences. This just goes to show that the consequences of an idea are irrelevant to the truth or reasonableness of an idea.

#### 4.3.6 Appeal to authority

In a society like ours, we have to rely on authorities to get on in life. For example, the things I believe about electrons are not things that I have ever verified for myself. Rather, I have to rely on the testimony and authority of physicists to tell me what electrons are like. Likewise, when there is something wrong with my car, I have to rely on a mechanic (since I lack that expertise) to tell me what is wrong with it. Such is modern life. So there is nothing wrong with needing to rely on authority figures in certain fields (people with the relevant expertise in that field)—it is inescapable. The problem comes when we invoke someone whose expertise is not relevant to the issue for which we are invoking it. For example, suppose that a group of doctors sign a petition to prohibit abortions, claiming that abortions are morally wrong. If Bob cites that fact that these doctors are against abortion, therefore abortion must be morally wrong, then Bob has committed the appeal to authority fallacy. The problem is that doctors are not authorities on what is morally right or wrong. Even if they are authorities on how the body works and how to perform certain procedures (such as abortion), it doesn't follow that they are authorities on *whether or not these procedures should be performed*—the *ethical* status of these procedures. It would be just as much an appeal to consequences fallacy if Melissa were to argue that since some other group of doctors supported abortion, that shows that it must be morally acceptable. In either case, since doctors are not authorities on moral issues, their opinions on a moral issue like abortion is irrelevant. In general, an appeal to authority fallacy occurs when someone takes what an individual says as evidence for some claim, when that individual has no particular expertise in the relevant domain (even if they do have expertise in some other, unrelated, domain).