

Finding, Clarifying, and Evaluating Arguments

Written by

[E.J. Coffman](#)

Associate Professor of Philosophy
[The University of Tennessee](#)

Edited by

[Trevor Hedberg](#)

Graduate Student
[The University of Tennessee](#)

Author's Note

In writing this essay, I've benefited tremendously from studying other similar works. I highly recommend the following three, which have been especially helpful to me:

- Richard Feldman's *Reason and Argument* (Prentice-Hall, 1999)
- Jim Pryor's "Philosophical Terms and Methods" (available at www.jimpryor.net/teaching/vocab/index.html)
- Eugene Mills' "The Nature of Argument" (unpublished).

Licensing and Distribution

This document is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivs 3.0 Unported License](#).



Table of Contents

I. Introduction	3
II. Finding Arguments.....	3
III. Clarifying an Argument	4
A. Identify the Argument’s Conclusion	4
B. Identify All Explicit Premises.....	4
C. Add Any Implicit Premises.....	4
D. Regiment the Argument.....	5
IV. Evaluating an Argument	6
A. Step One: Test the Argument’s Form	6
1. ‘Valid’ and ‘Invalid’ Defined	7
2. The Validity Test Explained	9
3. The Validity Test Illustrated	10
B. Test the Argument’s Premises.....	12
1. ‘Conditional Statement’ Defined	13
2. The Counterexample Method Explained	14
3. The Counterexample Method Illustrated	14
V. Summary	15
VI. Appendix 1: Some Common Argument Forms.....	16
Common Valid Argument Forms	16
Common Invalid Argument Forms	18
VII. Appendix 2: Analyzing Concepts	18

Seven Essential Points about Philosophical Analysis.....	19
Seven Essential Points about Philosophical Analysis: Breakdown	20
VIII. Appendix 3: Glossary of Key Terms	25

I. Introduction

In this course, we'll use the word 'argument' as it's typically used in academic settings. An **argument** is a set of statements or claims related to each other like this: one of the statements—the **conclusion**—is supposed to be rationally supported by the other statements—the **premises**. An argument is a line of reasoning offered in support of a particular claim or thesis.

One of the main goals of any philosophy course is to help you hone your ability to find, clarify, and evaluate arguments that others have presented, and to present and defend arguments of your own. In what follows, I'll offer some standard guidance on becoming better at finding, clarifying, and evaluating arguments about topics of all kinds. Along the way, I'll introduce some special terminology that philosophers use when discussing arguments (collected in [Appendix 3: Glossary of Key Terms](#)). Mastering this terminology will enable you to give especially clear, satisfying, and illuminating assessments of the arguments you encounter—assessments that pinpoint exactly what's wrong (and right) with an argument.

II. Finding Arguments

One of the best ways to become better at finding arguments is to actively look for them in whatever you happen to read, watch, or listen to. Here are a couple of written passages for practice:

[1] I have a few questions for those who have raised their voices against the recent Supreme Court decision to preserve our constitutional right to engage in symbolic acts of protest, including the burning of the American flag:

Are you as outraged when our Constitution is assaulted?

Did you protest when the constitutional rights of black citizens were denied? Did you work for their rights to vote, to equal education, to fair housing?

Have you spoken out against the assault on our Constitution by the illegal maneuverings of the boys in the White House during the Iran-Contra affair?

...In short, can you honestly say that you love your flag when you have been silent in protecting all that it stands for?¹

[2] Even the Fool is forced to agree that something-than-which-nothing-greater-can-be-thought exists in the mind, since he understands this when he hears it, and whatever is understood is in the mind. And surely that-than-which-a-greater-

¹ Vicki Lewin, Letter to the editor, *Rochester Times-Union*, July 12, 1989, p.7A (reprinted and discussed by Richard Feldman in his *Reason and Argument*).

cannot-be-thought cannot exist in the mind alone. For if it exists solely in the mind, it can be thought to exist in reality also, which is greater. If then that-than-which-a-greater-cannot-be-thought exists in the mind alone, this same that-than-which-a-greater-*cannot*-be-thought is that-than-which-a-greater-*can*-be thought. But this is obviously impossible. Therefore there is absolutely no doubt that something-than-which-a-greater-cannot-be-thought exists both in the mind and in reality.²

Does passage [1] express or contain an argument? How about passage [2]?

So far as I can see, passage [1] does not contain an argument, but passage [2] clearly does (indeed, [2] contains one of the most widely discussed arguments for God's existence in the history of philosophy). A bit of reflection on passage [2] reveals an important lesson: Arguments are not always expressed as clearly as they could be! Oftentimes, before we can properly *evaluate* an argument we've found, we'll have to do some work to *clarify* the argument—that is, to make the argument's **form** (its general logical structure, its underlying logical pattern) and **content** (its particular premises and conclusion) as clear as possible.

III. Clarifying an Argument

There's a standard method we can use to clarify an argument we've found (and want to properly evaluate). The method has four steps:

A. Identify the Argument's Conclusion

An argument's **conclusion** is its “main thesis,” the one statement that's supposed to be rationally supported by the other statements (the argument's **premises**).

B. Identify All Explicit Premises

An argument's **explicit premises** are the considerations (reasons, pieces of evidence) that the author *actually states* on behalf of the argument's conclusion.

C. Add Any Implicit Premises

An argument's **implicit premises** are any *unstated* considerations (reasons, pieces of evidence) that must be added to the argument's explicit premises so that the explicit premises are clearly connected to the argument's conclusion.

Some authors do a good job of laying all their premises on the table. When that happens, you won't have to spend a lot of time searching for implicit premises. Unfortunately, though, many authors—for whatever reason—fail to put all their premises in clear view. In that case, you may have to spend some time trying to figure out the implicit premises.

² St. Anselm, excerpt from *Proslogion*.

D. Regiment the Argument

To **regiment** an argument, first label each of its parts (its premises and its conclusion), then list each part on its own line. This exercise will not only clarify the argument's **form** and **content** but will also facilitate discussion and evaluation of the argument's **premises**.

Let me illustrate how to regiment an argument. Consider this passage:

Our universe must have been designed by an intelligent being. After all, our universe is extremely complex. And any extremely complex thing was designed by something intelligent.

Clearly, this passage contains an argument, a line of reasoning offered in support of a particular claim or thesis. To regiment that argument, we do two things: (1) Label each of the argument's parts, then (2) List each part on its own line. Completing these two tasks might yield something that looks like this:

P1: Every extremely complex thing was designed by an intelligent being.
P2: Our universe is extremely complex.
C: Thus, our universe was designed by an intelligent being.

In the process of regimenting the argument in the above passage, I tried to follow what philosophers often call the **Charity Principle**: When clarifying an argument, make the argument as sensible as you possibly can, given what its author said when presenting it. Adhering to the Charity Principle when trying to clarify an argument will help ensure that we don't waste precious time and energy considering an obviously hopeless argument. Here are two pieces of practical advice implied by the Charity Principle:

1. When clarifying an argument, eliminate any unnecessary premises and language.
2. When clarifying an argument, state and arrange the argument's premises so as to show how they are supposed to support the argument's conclusion. (This often involves making the argument's wording more uniform, and re-ordering the argument's premises and conclusion so that its underlying logical structure is easier to see.)

Now that we've regimented the argument in the above passage, the argument's **form**—its general logical structure, or underlying logical pattern—is much easier to see. Studying the regimented argument reveals that it has the following form:

P1: Every P is a Q.
P2: X is a P.
C: Thus, X is a Q.

In the next section, we'll talk more about how to figure out an argument's **form**.

Finally, notice that we could now use ‘P1’ to refer to the argument’s first premise, ‘P2’ to refer to the argument’s second premise, and ‘C’ to refer to the argument’s conclusion. Such labels illuminate an argument’s logical structure (by indicating which claims are premises and which claims are conclusions) and also facilitate discussion of the individual statements that make up an argument.

We’ve now seen how to *clarify* an argument. Once we’ve clarified an argument, we’re ready to *evaluate* it—that is, to figure out whether or not the argument’s premises really do constitute a good reason to believe its conclusion.

IV. Evaluating an Argument

There’s a standard method we can use to evaluate an argument that we’ve clarified via the techniques described in the last section. In order to evaluate an argument, we need to know what it takes for an argument to qualify as **good**. Basically, a **good argument** is one that has proper form *and* acceptable (plausible, reasonable) premises. On the other hand, a **bad argument** is one that has improper form *or* some unacceptable (implausible, unreasonable) premises.

So, to qualify as **good**, an argument must pass *two* tests: It must have proper form, and it must have acceptable premises. Naturally, then, the standard method for evaluating an argument involves *two* steps: the first step tests the argument’s **form**—its general logical structure or underlying logical pattern—while the second step tests the argument’s premises.

A. Step One: Test the Argument’s Form

When you’re testing an argument’s **form**, you’re trying to figure out the quality or strength of the logical connection between the argument’s premises and its conclusion. An argument that has *proper* logical form has a *strong* logical connection between its premises and its conclusion.

Recall this argument from the last section:

- P1: Every extremely complex thing was designed by an intelligent being.
- P2: Our universe is extremely complex.
- C: Thus, our universe was designed by an intelligent being.

Perhaps this argument won’t ultimately qualify as good, but it certainly has proper logical form or structure. For there’s a *very strong* logical connection between this argument’s premises and its conclusion: *If* this argument’s premises were true, its conclusion would *have* to be true too—the conclusion “falls right out” of the premises. So, this argument definitely passes the first test involved in our standard method for evaluating arguments: It has proper logical form.

Now, if we were to continue evaluating the above argument, we would at this point proceed to the *second* step of argument evaluation: We would ask ourselves whether each

of the argument's premises are acceptable (plausible, reasonable). And if we found each of those premises to be reasonable, we would then judge the argument to be good. On the other hand, if we found the argument to involve at least one dubious premise, we would judge the argument to be bad.

The following crucial fact about the *first* step of argument evaluation is now in view:

At the first step of argument evaluation, we *completely ignore* the question of whether the argument's premises are actually true. Instead, we focus *entirely* on the argument's form—its general structure or underlying pattern.

We now know that the *first* step of argument evaluation is to determine whether the argument has proper logical form. Interestingly, whether or not a particular argument is properly formed depends in part on what kind of argument it is. Basically, there are two kinds of arguments: **deductive arguments** and **inductive arguments**. A **deductive** argument aims to *prove* that its conclusion is true, to *establish* its conclusion beyond a shadow of a doubt. A mathematical proof is a good example of a deductive argument. An **inductive** argument, on the other hand, attempts only to show that its conclusion is *somewhat reasonable to believe*. The complex case a detective slowly builds for the guilt of a particular suspect is a good example of an inductive argument.

As it happens, most of the arguments we'll explore in this course are **deductive**. For the moment, then, we'll limit ourselves to discussing how to figure out whether a deductive argument that we've found and clarified is properly formed.

1. 'Valid' and 'Invalid' Defined

When a deductive argument is properly formed, we say that the argument is **valid**. Repeat: In this course, a **valid argument** is a deductive argument that has proper logical form, a properly formed deductive argument.

To qualify as **valid**, a deductive argument must be such that its conclusion would *have* to be true *if* its premises were true. To put it a slightly different way, a valid argument is one whose form is such that its conclusion *must* be true *if* its premises are true. To say it one last way, a valid argument is one whose form is truth-preserving in the sense that the form *always* yields a true conclusion *if* it is fed true premises.

Thinking about the following common **valid argument forms** should help you understand the concept of **validity** more deeply:

P1: Every P is a Q.

P2: X is a P.

C: X is a Q.

P1: If P, then Q.

P2: P.

C: Q. (This valid argument form is commonly called 'modus ponens'.)

P1: If P, then Q.
 P2: If Q, then R.
 C: If P, then R. (This valid form is commonly called ‘hypothetical syllogism’.)

P1: If P, then Q.
 P2: Not Q.
 C: Not P. (This valid form is commonly called ‘modus tollens’.)

Each of these argument forms is valid: Each is such that *if* it is fed true premises, it *must* yield a true conclusion. To put it a little differently, each is such that it *simply couldn’t* have all true premises and a false conclusion. In the next section, we’ll talk about how to figure out whether or not a particular argument form is valid.

It’s worth emphasizing that the facts about whether or not a particular argument is valid depend *solely* on the argument’s form, and thus are *completely independent of* the facts about whether the argument’s premises are *actually true*. As a result, you simply do not need to know whether an argument’s premises are actually true in order to figure out whether or not the argument is valid.

When a deductive argument has *improper* form, we say that the argument is **invalid**. An **invalid argument** is one whose form is *not* truth-preserving: an invalid argument form is such that *even if* it is fed true premises, it *could still* yield a false conclusion. Therefore, all invalid arguments are **bad**.

Thinking about the following common **invalid argument forms** should help you understand the concept of **invalidity** more deeply:

P1: Every P is a Q.
 P2: X is a Q.
 C: X is a P.

P1: If P, then Q.
 P2: Q.
 C: P. (This invalid form is commonly called ‘affirming the consequent’.)

P1: If P, then Q.
 P2: Not P.
 C: Not Q. (This invalid form is commonly called ‘denying the antecedent’.)

Each of these argument forms is invalid: Each is such that *even if* it is fed true premises, it *could still* yield a false conclusion. For example, to see that the second and third argument forms are invalid, plug *George W. Bush is on Mars today* into the “P” positions, and *George W. Bush is on a planet today* into the “Q” positions. **Exercise:** **Construct an argument that shows the first form to be invalid.**

In this course, we’ll spend a fair amount of time trying to figure out whether or not certain interesting and important philosophical arguments are valid. Fortunately, there’s

a good method we can use to test an argument for validity. We'll call this method the **Validity Test**.

2. The Validity Test Explained

The **Validity Test** involves two steps: (1) Figure out the argument's form, and (2) Try to prove that the form *isn't* truth-preserving.

a. Figure out the argument's form

To figure out an argument's form, we replace the argument's particular **content**—its particular terms, phrases, and clauses—with capital letters (P, Q, X, Y, ...) that will serve as general placeholders for the argument's particular content. Completing this exercise will yield a completely general **argument form**—something that looks like this:

P1: If P, then Q.
P2: P.
C: Q.

The exercise of replacing an argument's particular content with general placeholder letters may initially seem difficult; but this task becomes pretty easy once you've seen it done a few times. Let me illustrate how to figure out an argument's form by using an argument we considered earlier. Recall this one:

P1: Every extremely complex thing was designed by an intelligent being.
P2: Our universe is extremely complex.
C: Thus, our universe was designed by an intelligent being.

To isolate this argument's form, we replace its particular content with capital letters that will serve as general placeholders for the argument's specific content. Completing this replacement exercise will yield something that looks like this:

P1: Every P is a Q.
P2: X is a P.
C: Thus, X is a Q.

I used the capital letter 'P' as a placeholder for 'extremely complex thing'; 'Q' as a placeholder for 'thing designed by an intelligent being'; and 'X' as a placeholder for 'our universe'. Note also that I left 'every', 'is', and 'thus' intact, since these terms (unlike the ones I replaced) actually give the argument whatever general logical structure it has. The result is a completely general **argument form**.

We've now figured out the form exemplified by the above argument for the conclusion that our universe was designed by an intelligent being. Again, we did this by replacing the argument's particular content with capital letters that serve as general placeholders, while leaving intact all the terms that give the argument whatever logical structure it has.

Those who are initially confused about how to figure out an argument's form get the hang of it pretty quickly after seeing it done a few times. Here's another example:

- P1: If there's horrific evil in the world, then God doesn't exist.
 P2: If Buddhism is true, then God doesn't exist.
 C: If there's horrific evil in the world, then Buddhism is true.

To figure out this argument's form, replace its particular content with capital letters that will serve as general placeholders, leaving intact any terms that give the argument its general logical structure. Something like this will be the result:

- P1: If P, then Q.
 P2: If R, then Q.
 C: If P, then R.

I used 'P' as a placeholder for 'there's horrific evil in the world'; 'Q' as a placeholder for 'God doesn't exist'; and 'R' as a placeholder for 'Buddhism is true'. Notice also that I left the "If...then..." constructions alone, since those constructions actually give this argument its general logical structure. Again, the result is a completely general argument form, one that underlies the argument set out near the top of this page.

We've now seen how to figure out an argument's form. But how can we tell whether a particular argument form is valid or invalid?

b. Try to show that the argument form isn't truth-preserving

We've said that a valid argument form is truth-preserving—such that *if* it is fed true premises, it *must* yield a true conclusion. To determine whether a particular argument form is valid, we try to construct an argument *of that form* with *all true premises* and a *false conclusion*.

If you can construct such an argument, you'll have shown that the argument form under consideration is invalid—that is, *not* truth-preserving. Since all invalid arguments are bad, discerning that an argument has an invalid form shows that the argument is bad.

On the other hand, if you're not able to construct an argument of the relevant form with *all true premises* and a *false conclusion*, that's a good sign that the form under consideration is valid. You would then proceed to the second step of argument evaluation, where you carefully consider the argument's particular premises.

3. The Validity Test Illustrated

The **Validity Test** makes more sense once you've seen it done. So, let me illustrate the method by using it to test one of the arguments we encountered earlier. Recall this one:

- P1: If there's horrific evil in the world, then God doesn't exist.
 P2: If Buddhism is true, then God doesn't exist.
 C: If there's horrific evil in the world, then Buddhism is true.

Recall that the first step of the Validity Test is to *figure out the argument's form*. To do this, we replace the argument's particular content with general placeholder letters, while leaving all the "logically important" terms intact. When we replace the above argument's

particular content with general placeholder letters, we reveal an underlying argument form that looks like this:

P1: If P, then Q.
 P2: If R, then Q.
 C: If P, then R.

We now proceed to the second step of the Validity Test, where we try to construct a different argument *of the same form* with *all true premises* and a *false conclusion*. Constructing such an argument would show that the argument form under consideration is invalid, and so that the argument under consideration is bad. So, what do you think: is the form currently under consideration valid or invalid?

The argument form

P1: If P, then Q.
 P2: If R, then Q.
 C: If P, then R.

is invalid. We can prove this by constructing an argument *of that form* with *all true premises* and a *false conclusion*. Such arguments can be constructed; here's an example:

P1: If EJ is a human, then EJ is a mammal. [TRUE: *All humans are mammals.*]
 P2: If EJ is a dog, then EJ is a mammal. [TRUE: *All dogs are mammals.*]
 C: If EJ is a human, then EJ is a dog. [FALSE: *No human is also a dog.*]

Notice that this argument has the form currently under investigation (If P, then Q; If R then Q; So, if P then R). Notice also that the argument has *all true premises* but a *false conclusion*. The existence of this argument *proves* that the general argument form

P1: If P, then Q.
 P2: If R, then Q.
 C: If P, then R.

is invalid—that is, *not* truth-preserving.

Finally, notice that we've now shown that the argument we were *originally* concerned with—the one above about horrific evil, God, and Buddhism—is bad. For we have shown that the original argument's underlying form (If P, then Q; If R then Q; So, if P then R) is invalid, and as you'll recall, every argument with an invalid form is bad.

I've now illustrated the standard way to test an argument for validity—the Validity Test. To perform the Validity Test on an argument, (1) Figure out the argument's form, then (2) Try to construct a different argument *of the same form* with *all true premises* and a *false conclusion*. If you succeed, you'll have *proved* that the original argument is improperly formed, and therefore bad. On the other hand, if don't succeed, that's a good sign that the argument whose form you're testing is valid. You'd then proceed to the

second and final step of argument evaluation, where you scrutinize the argument's premises.

B. Test the Argument's Premises

To test an argument's premises, ask yourself this question: *Is it reasonable for me to believe all the premises involved in this argument?*

If your answer is a resounding 'Yes', then the argument qualifies as **good**: it is valid *and* has acceptable (plausible, reasonable) premises. On the other hand, if you have doubts about one or more of the argument's premises, then the argument is **bad** (at least until something has been done to mitigate your doubts about its premises).

Let me briefly highlight a thought that the previous paragraph may prompt in some readers. There's an important—though ultimately harmless—way in which the concepts of **good argument** and **bad argument** are “subjective” or “person-relative”. It's possible that different people properly arrive at different verdicts on the question whether a particular argument is good: One person may properly evaluate an argument as *good*, while another person properly evaluates the same argument as *bad*. How can this be?

What makes for this possibility is the fact that people can differ with respect to the evidence they possess. For example, professional scientists have large bodies of evidence for initially odd-sounding scientific theories, evidence that non-scientists lack. Because people can differ with respect to the evidence they possess, people can differ with respect to what's reasonable for them to believe: What's reasonable for you to believe is determined largely (if not completely) by your evidence.³ For example, professional physicists have evidence from various scientific experiments that makes it reasonable for them to believe that light sometimes behaves like a wave but at other times behaves like a particle. But it would be *unreasonable* for a typical person on the street who *lacks* the relevant scientific evidence to believe that light has those apparently incompatible features.

So, because people can differ with respect to the evidence they have, people can differ with respect to what's reasonable for them to believe. And this makes for the possibility that different people properly arrive at different verdicts on the question whether a given argument is good. One way people can try to resolve such disagreement is to share with each other all the (shareable) evidence they have that's relevant to the argument under consideration.

³ And this, incidentally, is a completely “objective” matter—that is, your evidence supports some claims and doesn't support others, and these facts are unaffected by your beliefs about what your evidence does or doesn't support. Similarly, the facts about what's reasonable for you to believe are objective: some things are reasonable for you to believe while other things aren't, and these facts are unaffected by your beliefs about what's reasonable for you to believe. For example, my simply *thinking* or *believing* that it's reasonable for me to believe that I'm a good athlete doesn't *automatically* make it reasonable for me to believe that I'm a good athlete.

Let's return to our main question: How do we test an argument's premises? As it happens, many of the arguments we'll encounter in this course involve premises that are **conditional statements** (for short, **conditionals**). *Roughly*, a **conditional** is a statement or claim of the form "If P, then Q." To evaluate an argument involving conditionals, you'll need to know how to figure out whether a particular conditional is reasonable for you to believe. Fortunately, there's a good method we can use to evaluate conditionals: the **Counterexample Method**. I'll describe and illustrate this method in a moment.

First, though, we need to say a bit more about conditionals. I just said that a conditional is, *roughly*, a statement of the form "If P, then Q." For the purposes of this course, we'll need a slightly more precise definition.

1. 'Conditional Statement' Defined

For our purposes, a **conditional statement** is any statement that *has the same meaning as* a statement of the form "If P, then Q." Interestingly, there are many different ways to say "If P, then Q." All these different ways of saying "If P, then Q" qualify as conditional statements. Here are the main kinds of conditionals we'll encounter in this course:

- If P, then Q. ("If I went to the store, then I went somewhere.")
- P suffices for Q. ("My going to the store suffices for my going somewhere.")
- P is a sufficient condition for Q.
- P implies Q.
- P entails Q.
- P only if Q.
- P requires Q.
- Q is required for P.
- Q is necessary for P.
- Q is a necessary condition for P.

Repeat: For our purposes, a **conditional statement** is a statement that *has the same meaning as* a statement of the form "If P, then Q." Since each of the statements just listed has the same meaning as a statement of the form "If P, then Q," each of those statements qualifies as a **conditional**.

Finally, in a statement of the form "If P, then Q," the statement in the "P" position is called the **antecedent**, and the statement in the "Q" position is called the **consequent**. For example, in the conditional

"If I went to the store, then I went somewhere"

I went to the store is the **antecedent**, and *I went somewhere* is the **consequent**.

We now know what a conditional is (a statement equivalent to a claim of the form "If P, then Q"), and what its parts are called (antecedent, consequent). We're finally ready to learn the standard method for evaluating conditionals: the **Counterexample Method**.

2. The Counterexample Method Explained

To use the **Counterexample Method** to evaluate a conditional statement, first translate the statement into “If P, then Q” form, and then try to imagine a coherent scenario in which the conditional’s antecedent holds but its consequent *doesn’t*.

If you succeed in imagining a coherent scenario in which the conditional’s antecedent is *true* but its consequent is *false*, you have found a **counterexample** to that particular conditional, thereby showing the conditional to be false.

If the Counterexample Method seems somewhat mysterious at first, think of it this way. A conditional says that *whenever* its antecedent holds, its consequent will *also* hold: a conditional says, in effect, that there *couldn’t* be a situation in which its antecedent is *true* but its consequent is *false*. Thus, describing a coherent scenario in which the conditional’s antecedent is true but its consequent is false shows that the conditional itself is false.

3. The Counterexample Method Illustrated

Like many other things we’ve discussed, the Counterexample Method for evaluating conditionals makes more sense once you’ve seen it used a few times. So, let me illustrate the method by using it to evaluate this conditional statement:

Firmly believing a claim that is in fact true suffices for your *knowing* the claim to be true.

To evaluate this statement with the Counterexample Method, we first translate it into “If P, then Q” form, and then try to find a counterexample to it. We now know—perhaps by consulting [the list on p.13](#)—that the statement we want to evaluate is equivalent to the following statement:

If you firmly believe a claim that is in fact true, then you *know* the claim to be true.

We have now completed the first step of the Counterexample Method: We’ve translated the statement we want to evaluate into “If P, then Q” form. We can now proceed to the second step of the Counterexample Method, where we try to imagine a coherent scenario in which the statement’s antecedent holds but its consequent doesn’t.

As it happens, such scenarios exist. Consider this one:

You desperately need fifty bucks. As a result, you (somewhat foolishly) bet a friend fifty bucks that the next time she flips a coin, it will land heads. Your friend accepts the bet, and promptly flips and catches a coin. Neither of you has seen the coin yet. Due entirely to “wishful thinking,” however, you *firmly* believe that the coin landed heads. And as it happens, the coin *did* land heads. Still, you do not yet *know* that the coin landed heads.

This imaginary scenario is a counterexample to the conditional now under consideration. The scenario is perfectly coherent. Further, it is one in which the conditional's antecedent ("You firmly believe a claim that is in fact true") holds but the conditional's consequent ("You *know* the claim to be true") *doesn't*. Recall that, in the scenario, you firmly believe the true claim that the coin landed heads, but you do not yet *know* that the coin landed heads. The scenario thus shows the conditional under consideration to be false. (Incidentally, we find counterexamples to the alleged sufficiency of true belief for knowledge as far back as Plato's dialogue *Theaetetus*.)

Two final notes about the **Counterexample Method**. First, the exercise of trying to find a counterexample to a particular conditional is often called **conducting a thought experiment**. Second, to qualify as a successful counterexample, an imagined scenario needn't be mundane or ordinary or familiar. Indeed, a counterexample may involve quite far-fetched circumstances, *so long as the example remains coherent*. And it should here be said that a large part of what makes Philosophy so fun yet so challenging is that it's often very hard to say whether a particular imaginary scenario is indeed *coherent*.

V. Summary

I've offered some standard guidance about how to strengthen your ability to *find*, *clarify*, and *evaluate arguments*. Improving your ability to do these things is a central goal of any philosophy course, and these skills are well worth developing.

There's no standard method for *finding* arguments. To become better at finding arguments, reflect on what an argument is (a line of reasoning offered in support of a particular claim or thesis), and then start looking for them in whatever you happen to read, watch, or listen to.

There are standard methods we can use to *clarify* and *evaluate* arguments we've identified. Frequently, in order to properly evaluate an argument you've found, you'll first need to clarify the argument. To clarify an argument, follow these four steps:

1. Identify the argument's **conclusion**.
2. Identify all **explicit premises**.
3. Add any **implicit premises**.
4. **Regiment** the argument (which involves following the **Charity Principle**).

Following these steps will typically yield an interesting argument whose **form** and **content** are quite clear—just the kind of argument that lends itself to proper evaluation.

To *evaluate* an argument, follow these two steps:

1. Perform the **Validity Test**.
2. Determine whether *all* the argument's premises are reasonable.

A **valid** argument is one whose underlying form is truth-preserving. To perform the **Validity Test** on an argument,

- a. figure out the argument's form, and then
- b. try to construct a different argument of the same form with all true premises but a false conclusion.

If you succeed in constructing such an argument, you'll have *proved* that the original argument has an **invalid** form and is therefore a **bad** argument. On the other hand, if you don't succeed, that's a good sign that the original argument is valid. Having completed the first step of your evaluation, you can proceed to the second step, which concerns the argument's particular content.

Many arguments involve **conditional statements**. To evaluate a **conditional**, use the **Counterexample Method**:

- a. translate the statement into "If P, then Q" form, and then
- b. try to imagine a **counterexample**—a coherent scenario in which the conditional's **antecedent** is *true* but its **consequent** is *false*.

If you succeed in imagining such a scenario, you'll have shown the conditional under consideration to be false. On the other hand, if you don't succeed, that's a good sign that the conditional is true.

Finally, if an argument is valid *and* has all reasonable premises, then it qualifies as **good**. But if an argument is invalid *or* has some questionable premises, then it is **bad**. Simply put, a **good argument** is a **valid argument** all of whose premises are reasonable.

VI. Appendix 1: Some Common Argument Forms⁴

Here are some common **valid** and **invalid argument forms**. Most of the arguments we encounter this semester will exemplify one or another of the following **forms**. Coming to *see* that each of the **valid forms** is in fact **valid** (truth-preserving)—and that each of the **invalid forms** is in fact **invalid** (*not* truth-preserving)—will help you achieve a better understanding of these crucial concepts.

Common Valid Argument Forms

- Affirming the antecedent (*modus ponens*)

P1: If P, then Q.

P2: P.

C: Q.

⁴ This section draws on Richard Feldman's helpful collection of common argument forms in his *Reason and Argument*.

- Denying the consequent (*modus tollens*)

P1: If P, then Q.

P2: Not Q.

C: Not P.

- Argument by elimination (I)

P1: Either P or Q.

P2: Not P.

C: Q.

- Argument by elimination (II)

P1: Either P or Q.

P2: Not Q.

C: P.

- Hypothetical syllogism

P1: If P, then Q.

P2: If Q, then R.

C: If P, then R.

- Simplification (I)

P1: P and Q.

C: P.

- Simplification (II)

P1: P and Q.

C: Q.

- Contraposition

P1: If P, then Q.

C: If Not-Q, then Not-P.

The following additional valid argument forms don't have standard names:

P1: P if and only if Q.

P2: Not P.

C: Not Q.

P1: All A-s are B-s.

P2: X is an A.

C: X is a B.

P1: All A-s are B-s.
P2: X is not a B.
C: X is not an A.

P1: All A-s are B-s.
P2: All B-s are C-s.
C: All A-s are C-s.

P1: No A-s are B-s.
P2: X is an A.
C: X is not a B.

Common Invalid Argument Forms

- Denying the antecedent

P1: If P, then Q.
P2: Not P.
C: Not Q.

- Affirming the consequent

P1: If P, then Q.
P2: Q.
C: P.

Unlike the first two **invalid argument forms**, the following two don't have standard names:

P1: All A-s are B-s.
P2: X is not an A.
C: X is not a B.

P1: All A-s are B-s.
P2: X is a B.
C: X is an A.

VII. Appendix 2: Analyzing Concepts⁵

Like Physics, Chemistry, and Biology, Philosophy is an academic discipline. But while you probably have a pretty good idea of what physicists, chemists, and biologists study, you may still be wondering what *philosophers* study. What's the subject matter of *Philosophy*?

⁵ For a complementary and very helpful discussion of analyzing concepts, see the relevant section of Jim Pryor's "[Philosophical Terms and Methods](#)."

To a large extent, the subject matter of Philosophy is certain abstract yet deeply interesting and important *concepts*—examples include *knowledge*, *freedom*, *causation*, *rightness*, and *justice*. Like physicists, chemists, and biologists, philosophers aim to shed light on the subject matter of their discipline. More precisely, philosophers aim to provide **good analyses** of interesting and important abstract concepts like those just mentioned.

In this course, we'll encounter a number of proposed analyses of some of our most interesting concepts, and we'll work to understand and properly evaluate those analyses. To understand and properly evaluate an analysis of a particular concept, you'll need to know the answers to three crucial questions:

1. What *is* a philosophical analysis?
2. How do contemporary philosophers present their proposed analyses?
3. What makes for a *good* analysis?

In my experience, when it comes to learning how philosophers analyze concepts, some people prefer to just hit the highlights, while others want to know all the details behind the highlights. In what follows, I'll try to please both kinds of people. For those who would prefer to just hit the highlights, I've boiled the facts about philosophical analysis down to seven essential points, which I'll set out in a moment. For those who would like to know the details behind the highlights, I'll restate the seven essential points, explaining each one and tying them all together.

Seven Essential Points about Philosophical Analysis

- [1] A philosophical analysis is a kind of definition.
- [2] A definition says that one word or group of words is *equivalent in meaning to* another word or group of words. In short, a definition is a claim of equivalence.
- [3] Thus, from [1] and [2], we can see that a philosophical analysis is an equivalence claim: It says that a certain interesting complex concept *is equivalent to* a certain combination of simpler (better understood, more intelligible, more familiar, clearer) concepts.
- [4] The sentence form “P if and only if Q” is a convenient and precise way to make an equivalence claim.
- [5] Contemporary philosophers often present their analyses using the “P if and only if Q” format. In this format, the concept being analyzed—the “target concept”—appears on the left side (where the ‘P’ is), while the concepts used to analyze the target concept—the “defining concepts”—appear on the right side (where the ‘Q’ is).
- [6] To qualify as *good*, an analysis must have the following two features: (1) Its two sides must really be equivalent (i.e., in *any* scenario where one side holds, the

other side holds too), *and* (2) The defining concepts must be more intelligible (clearer, simpler, more familiar) than the target concept.

- [7] To determine whether the two sides of a proposed analysis really are equivalent, you try to imagine an example in which one side holds but the other *doesn't*. If you can imagine such a scenario, the analysis is bad. If you *can't* imagine such a case, then the two sides of the analysis may well be equivalent, and the analysis may well be good.

Seven Essential Points about Philosophical Analysis: Breakdown

For those who would like to know the details behind the above highlights, I'll repeat the seven points above, along the way explaining each point and tying them all together.

[1] A philosophical analysis is a kind of definition.

When you offer an *analysis* of a concept, your goal is to provide a *good definition* of that concept, one that uses more familiar or better understood terms or concepts. Providing such a definition would shed some light on the concept in question: such a definition would make the target concept clearer and more intelligible.

[2] A definition says that one word or group of words is *equivalent in meaning to* another word or group of words. In short, a definition is a claim of equivalence.

To see that this is so, consider three related definitions:

Rectangle =	a parallelogram with four right angles
Parallelogram =	a quadrilateral whose opposite sides are both parallel and equal in length
Quadrilateral =	a closed shape with four straight sides

Each of these definitions highlights a particular word, then tells us that the word is equivalent in meaning to (i.e., has the same meaning as) another group of words. For example, the first definition highlights the word 'rectangle' and tells us that this word means the same thing as 'parallelogram with four right angles'.

Combining points [1] and [2] yields point [3].

[3] So, a philosophical analysis is an equivalence claim: it says that a certain interesting complex concept *is equivalent to* a certain combination of simpler (better understood, more intelligible, more familiar, clearer) concepts.

We now know what a philosophical analysis *is*: it's an attempt to shed light on an interesting complex concept by defining the concept in terms of simpler, clearer concepts. The next thing we need to know is how contemporary philosophers express or present their proposed analyses. [4] is our first step in that direction.

[4] The sentence form “P if and only if Q” is a convenient and precise way to make an equivalence claim.

To begin to see that this is so, think back to our [earlier discussion of conditional statements](#). You’ll recall that a statement of the form “If P, then Q” has the same meaning as a statement of the form “P is a sufficient condition for Q.” That is, “If P, then Q” can be rewritten as “P is a sufficient condition for Q.” You may also recall that “P only if Q” has the same meaning as “Q is a necessary condition for P.” That is, “P only if Q” can be rewritten as “Q is a necessary condition for P.”

We’re now in a position to see a crucial fact about the sentence form “P if and only if Q”: “P if and only if Q” is a convenient way to say *two very different things*. The first is this:

P if Q = If Q then P = Q only if P = Q requires P = P is necessary for Q

(If you don’t see right away that all these sentences are equivalent, [review our discussion of conditionals](#).)

Here’s the *second* thing “P if and only if Q” says:

P only if Q = P requires Q = If P then Q = P suffices for Q = P is sufficient for Q

What we’ve seen so far is this: “P if and only if Q” says that P is both a necessary *and* a sufficient condition for Q. “P if and only if Q” is a convenient way to say “P is a necessary and sufficient condition for Q”.

Now, to fully understand why point [4] holds, you need to see that “P is a necessary and sufficient condition for Q” is just a precise way to say “P is equivalent to Q”. Consider: If one condition is both necessary *and* sufficient for another condition, then those two conditions are *equivalent* in that they hold in *exactly* the same situations—you *can’t* have one condition without the other.

Take, for a simple example, the conditions of *being a bachelor* and *being an unmarried man*. *Being a bachelor* is a necessary and sufficient condition for *being an unmarried man*. (If you don’t see this immediately, you will after thinking about it a bit.) So, *being a bachelor* and *being an unmarried man* are equivalent: the two conditions hold in *exactly* the same circumstances—there *can’t* be an object that meets one of these conditions without meeting the other.

To sum up, then: “P if and only if Q” is a convenient way to say “P is a necessary and sufficient condition for Q”. Further, “P is a necessary and sufficient condition for Q” is a precise way to say “P is equivalent to Q”. We’ve now arrived at point [4]: the sentence form “P if and only if Q” is a convenient and precise way to make an equivalence claim.

In light of [4], [5] should make some sense.

[5] **Contemporary philosophers often present their analyses using the “P if and only if Q” format. In this format, the concept being analyzed—the “target concept”—appears on the left side (where the “P” is), while the concepts used to analyze the target concept—the “defining concepts”—appear on the right side (where the “Q” is).**

Contemporary academic philosophers want to use the best available tools or methods to do their work. One of the best available tools or methods for presenting an analysis is the “P if and only if Q” format. So, it makes sense that contemporary academic philosophers often use that format when stating their proposed analyses.

To illustrate, here’s a simple analysis of the concept of **knowledge**, written three different ways:

- You *know* a particular proposition (claim, statement, thesis) P **if and only if** you believe P to be true *and* P is in fact true.
- You *know* a particular proposition P **iff** you believe P *and* P is in fact true.
- You *know* a particular proposition P = **df.** You believe P *and* P is true.

All three of these sentences have the same meaning—they are just three different ways to say the same thing. In particular, the unfamiliar expressions “iff” and “= df.” are just different abbreviations of the expression “if and only if.” Since you now understand what the first sentence says *and* why it’s a handy way to present a philosophical analysis, you’ll automatically understand the second and third sentences too.

It’s worth noting that philosophers writing prior to the twentieth century didn’t use these formats to present their analyses. That’s not because those authors didn’t present analyses: they certainly did. It’s just that previous authors had different—often less convenient and precise—ways of expressing their analyses. (Philosophical texts from ancient Greece down to the present day contain claims that can be translated into the above formats.)

We now have some grasp on what an analysis *is* and how philosophers typically express their proposed analyses. The last thing we need to know is how to *evaluate* a proposed analysis. What makes an analysis *good*?

[6] **A good analysis of a particular concept must have the following two features: (1) Its two sides must really be equivalent, *and* (2) The “defining concepts” must be more intelligible (clearer, more familiar) than the “target concept.”**

As we’ve already seen, like *any* definition, a philosophical analysis will have *two* parts: (1) a part that highlights the concept to be analyzed (defined, clarified, explained) and (2) a part that tries to analyze (define, clarify, explain) the target concept.

In the three standard formats displayed above, the left side is the part that highlights the concept to be analyzed, while the right side is the part that tries to analyze the target concept. Recall, for instance, this sample analysis of **knowledge**:

You *know* a particular proposition $P = \mathbf{df.}$ You believe P *and* P is true.

The left side highlights the concept to be analyzed (in this case, propositional knowledge) whereas the right side offers an analysis or definition of the target concept. According to this analysis, *knowledge* is equivalent to *true belief*.

To learn how to evaluate a proposed analysis, we'll need to remind ourselves of the purpose of giving an analysis. What are you trying to do when you give an analysis? You're trying to shed light on a somewhat unclear and complex though interesting concept by successfully defining it in terms of simpler, clearer, better understood concepts. This suggests that to qualify as **good**, an analysis must meet two conditions:

- (1) Its two sides must really be equivalent, in the sense that we can't imagine them "coming apart": in *any* conceivable example (case, situation, scenario), *either* both conditions hold *or* neither condition holds.
- (2) The concepts used to analyze the target concept must be clearer (more intelligible, better understood) than the target concept.

So, a proposed analysis can fail in one of two ways: (1) Its sides really aren't equivalent, or (2) The "defining concepts" are no clearer than the "target concept". (As we'll see in this course, some proposed analyses fail *both* of these tests.) Discerning whether the defining concepts are clearer than the target concept is (relatively) straightforward. What needs some discussion is how we figure out whether the sides of a given analysis really are equivalent. That brings us to the last of our seven essential points.

[7] To determine whether the two sides of a proposed analysis really are equivalent, you try to imagine an example in which one side holds but the other *doesn't*. If you can imagine such a scenario, the analysis is bad. If you *can't* imagine such a case, then the two sides of the analysis may well be equivalent, and the analysis may well be good.

To determine whether the two sides of a given analysis really are equivalent, you try to come up with what we earlier called a counterexample to the analysis. That is, you try to imagine or describe a coherent example (case, situation, scenario) which shows that the two sides can "come apart," an example in which one side holds but the other *doesn't*. If you come up with such a case, you'll have shown that the two sides of the analysis really *aren't* equivalent, and so that the analysis fails. (Remember: An analysis is correct only if its two sides really are equivalent—only if the concepts on the right side involve neither *more* nor *less* than the concept being analyzed.) On the other hand, if you *can't* come up with a

case in which the two sides of a proposed analysis “come apart,” that’s a good sign that the sides really are equivalent.

VIII. Appendix 3: Glossary of Key Terms

- Analysis**— a type of equivalence claim—a statement that a certain interesting complex concept is equivalent to a certain combination of simpler concepts
- Bad**— An analysis is bad if *either* its two sides are not equivalent *or* the defining concepts are not more intelligible (clearer, simpler, more familiar) than the target concept.
- Good**— An analysis is good if its two sides are equivalent *and* the defining concepts are more intelligible (clearer, simpler, more familiar) than the target concept.
- Antecedent**— In a conditional statement, the phrase or clause that corresponds to the “P” position
- Argument**— a line of reasoning offered in support of a particular claim or thesis; more precisely, a set of statements in which a conclusion is supposed to be rationally supported by premises
- Bad**— An argument is bad if it has *either* improper form *or* unacceptable premises.
- Deductive**— A deductive argument aims to prove that its conclusion is true.
- Good**— An argument is good if it has proper form *and* acceptable premises.
- Inductive**— An inductive argument attempts to show that its conclusion is at least somewhat reasonable to believe (as opposed to decisively establishing its conclusion).
- Invalid**— An argument is invalid if its form is such that its conclusion could be false even if its premises are true.
- Valid**— An argument is valid if its form is such that its conclusion must be true if its premises are true.
- Charity Principle**— When clarifying an argument, make the argument as sensible (i.e., as likely to succeed) as you possibly can, given what its author said when presenting it.
- Conclusion**— the statement in an argument that’s supposed to be supported by the remaining statements
- Conditional Statement**— any statement that has the same meaning as a statement of the form “If P, then Q”

Consequent—In a conditional statement, the phrase or clause that corresponds to the “Q” position

Counterexample (to a Conditional Statement)— a coherent scenario in which a conditional’s antecedent is true but its consequent is false

Counterexample Method— To use the **Counterexample Method** to evaluate a conditional statement, first translate the statement into “If P, then Q” form, and then try to imagine a coherent scenario in which the conditional’s antecedent holds but its consequent doesn’t.

Form (of an Argument)— the argument’s general logical structure or underlying logical pattern

Premises— the statements in an argument that are supposed to support the remaining statement

Explicit— considerations (reasons, pieces of evidence) that the author actually states on behalf of the argument’s conclusion

Implicit— any unstated considerations (reasons, pieces of evidence) that must be added to the argument’s explicit premises so that the explicit premises are clearly connected to the argument’s conclusion

Regiment an Argument— To regiment an argument, label each of its parts (i.e., its premises and conclusion), and list each part on its own line.

Thought Experiment— imaginary scenario that tests the acceptability of a particular conditional statement

Validity Test— To test the validity of an argument, figure out the argument’s form, and then try to prove that the form *isn’t* truth-preserving.