

2. Critical Thinking I: How to be a Reasonable Person

Critical thinking is about reasoning clearly and effectively to achieve our goals. The traditional goal in critical thinking courses is getting at the truth and avoiding falsehoods. Our conception of critical thinking can be expanded to include other goals and appropriate standards for achieving these. So, for instance, an architect isn't concerned with getting at the truth, but with good design. So, an architect will be concerned with standards and methods appropriate to achieving good design. But even in this context, it would be irrational for an architect to ignore truths about, say, the load bearing properties of various building materials. Whatever our goals, getting at the truth about many things will matter. So, truth oriented critical thinking in the context of inquiry is a vital and indispensable foundation for reasoning effectively.

Biases and confusions about the nature of truth and reason are widespread. So, we are going to start here by introducing the metaphysical and conceptual framework typically presupposed in inquiry. We will also discuss the personal traits and social conditions that are required for inquiry to proceed productively. In the next chapter we will get acquainted with the methods and skills employed in reasoning clearly and effectively. This will include a basic introduction to logic and logical fallacies.

Subjects and Objects

Let's start with the modest metaphysical assumption that we all live on planet Earth. This means we have a shared reality. One that is populated with various and sundry objects (or better, containing stuff that can be divided in to objects in any number of ways). This is the realm of objects, or the objective world. As embodied creatures, we are among the objects populating the objective realm. But in addition to being objects we are also subjects.

As subjects we have some experience of our shared reality, the objective realm. But, our experience of the world is limited by our perspectives. Further our impressions and beliefs are liable to be distorted by biases and assorted other intellectual bad habits. So, one thing we can all recognize about being subjects is that our impressions, beliefs and opinions are fallible. We are limited and imperfect in ways that make error quite possible. That's just life as a subject, having subjective impressions and beliefs means being prone to error. Fortunately, we can always expand the basis of evidence we reason from by sharing our impressions and beliefs with each other. And we can improve the reliability of our reasoning by cultivating the intellectual habits that steer us away from biased and distorted ways of thinking. This is the point of critical thinking.

Philosophers typically use the word “subjective” to refer to what is mind dependent. All of your thoughts, feelings and beliefs exist in the subjective realm in this sense of the word. But then your beliefs are about the objective world. When I believe I have chocolate ice cream in the freezer, I’m representing a part of the objective world as being a certain way. So, my belief while it exists in the subjective realm as a property of me, a subject, is about the objective realm and it can accurately represent an aspect of the objective world. In which case my belief is true. Or it can fail to fit the way things are, say, if my son has finished the chocolate ice cream. In this case my belief is false.

So, your beliefs, are aspects of you, a subject, but they aim at representing things that are going on in our shared reality, aspects of the objective realm. For your belief to be true is just for your representation of how things are in the objective realm to fit or correspond with what is actually going on in the objective realm. That is, your belief is true when it represents some aspect of our shared reality the way it is.

Subjective realm includes whatever depends on your mind as a subject	Objective realm includes all aspects of our shared reality
All your thoughts, sense impressions, feelings, beliefs, fears and hopes are aspects of you as a subject.	The objective realm includes all the things, states of affairs and ways things are independent of you as a subject.

As subjects, it is generally good for us to have true beliefs and avoid false beliefs. When we have true beliefs, we are more capable of acting effectively, achieving goals, avoiding hazards, and generally having a good time. I suppose this is a value statement, but not the sort of value statement anyone is likely to dispute. This much of the value of having true beliefs comes along with being subjects who have needs and goals in a world full of objects (and subjects) that can be helpful or harmful to us. So, special cases aside, it’s good to have true beliefs.

Truth

We just claimed that for your belief to be true is for it to represent things as they are. The basic idea here is that beliefs and claims are true when they correspond to how things are in our shared reality. Ordinary declarative sentences say something about how things are. What a declarative sentence says represents some aspect of reality. So, take a few everyday examples:

- There is a spruce tree in Stuart’s front yard.
- Lake Washington is east of downtown Seattle
- Your keys are on the kitchen counter.

Each of these sentences represents some aspect of our shared reality as being a certain way. The sentence is true if that part of the world is the way the sentence says it is. Since truth is about correspondence with our shared reality, it concerns the objective realm. Truth is objective.

Our society is currently riddled with some confusing ways of talking about truth. We have become accustomed to talking about “my truth” or “your truth.” But if we stay focused on our ordinary understanding of truth as correspondence to reality, we can avoid confusions about truth being subjective or relative. Corresponding to our shared reality is obviously what we mean when we count the belief that Lake Washington is east of Seattle as true. For my belief or claim to be true is just for it to represent some aspect of our shared reality as it is. What is true depends only on how things are objectively in our shared reality. Beyond shaping what is true about my own mind, I can’t make things true merely by willing, wishing or believing them.

Given this ordinary everyday understanding of truth, it should be clear the truth doesn’t belong to anyone. Nobody gets to dictate, define or decide what is the case, except in the very limited respect where a person decides what to do, how to think, or who to be. As a subject, I have this much power to shape our shared reality and no more. There is no “my truth” or “your truth.” The only way to make sense out of truth being subjective would be to deny the existence of a shared reality. Truth could be subjective only if I live in my own little world and you live in your own separate reality. This would be to reject the very modest metaphysical assumptions we started this chapter with. I can’t prove that we have a shared reality, but not having one sure sounds lonely.

When I believe something, I take it to be true. I suspect this is all most talk of “my truth” or “your truth” amounts to, a confusing way of talking about what we believe. But this kind of talk involves a rhetorical cheat in suggesting that my belief, which could well be just plain false, is still somehow to be associated with what is true. Worse, talk of “my truth” and “your truth” blurs the difference between the subjective and the objective.

Finally, before we leave the topic of truth, let’s consider the difference between these two questions:

- What is it for a claim to be true?
- How do we determine that a claim is true?

It’s important to keep these two questions separate. Questions about how we know whether something is true are epistemic questions. These questions are concerned with how our minds relate to the world. But the question of what it is for something to be true is not an epistemic issue. The truth of a claim is quite independent of how or whether we know it to be true. There are many truths we don’t know and some of the things we think we know just aren’t true. If you are not sure about this, consider these two claims:

- There is intelligent life on other planets.
- There is no intelligent life on other planets.

One of these claims is true. We can be sure of this on the basis of logic alone. Either claim being false would make the other true. We don't know which of these two claims is true and yet one of them is true. Whichever of these claims is true, its being true doesn't depend on whether we know it to be true. There are many truths that will never be known or believed by anyone, and appreciating this is enough to see that the truth of a claim is not relative to belief, knowledge, proof, or any other epistemic notion.

So, what it takes for a claim to be true doesn't depend on what we believe, or what we think we know, (except in the special case of claims about what we believe). What it takes for a claim to be true only depends on what on how things are in our shared reality. Once we get clear on subjects, objects and truth, the answer to our first question above is pretty clear. All it is for a claim to be true is for what it says to fit with how things are. But, naturally most of us are more concerned with how we can determine when the claims people make and the things we believe are true. This question is more challenging. It's also what critical thinking and epistemology are about. We will get to the issue of how to determine whether a claim is true or false in the next chapter when we turn to the basic methods for evaluating reasons and evidence.

Rationality

For your beliefs to be rational, or reasonable (we'll treat these terms as synonyms), is just for them to be held on the basis of the best available reasons. To be reasonable, in the literal sense of the word, is to be amenable to reason. That is, the reasonable person is the person who forms or revises beliefs by **yielding** to the best reasons.

Good reasons are reasons that are truth-oriented. So, all it means for your belief to be rational, or reasonable, is for it to be held for the most truth-oriented reasons available to you. This much should make it clear why it is good to be rational. Being rational is more likely to get you true beliefs and true beliefs are good because they help you act effectively, achieve your goals, avoid hazards, and they give you a shared basis for understanding and communicating with others.

[We should note that the words "rational" and "reasonable" can also refer to choosing or acting in ways that aim at maximizing some goal or value. Words are often ambiguous. The way to be comfortable with ambiguity is to get clear on how words are being used and to track the various usages. Talk of rational or reasonable **belief** (as opposed to *choice* or *action*) can generally be understood as truth-oriented simply because to believe something is to take it to be true].

Rationality is not a kind of human imposed authority over what is true or what we should believe. The only thing that is authoritative concerning what we should believe is how

things are in our shared reality. Again, to believe something is to take it to be true. To believe rationally is just to believe in a ways that target the truth well. To believe irrationally is to aim badly at the truth. Rational belief isn't guaranteed to hit the target of truth. But irrational belief involves a kind of unforced error.

Talk of rationality, objectivity and truth have some difficult connotations in the minds of many. These concepts often get associated with things like maleness, authority, power, and now even "whiteness." These difficult associations appear to be based on antiquated stereotypes of one sort or another. But thinking based on stereotypes is highly unreliable and perhaps we are in a position now to see how thinking in terms of stereotypes misleads many of us concerning these fundamental concepts. To be reasonable literally means to be amenable to good reasons (this especially includes the good reasons of others who think differently). To be a reasonable or rational believer involves a good measure of intellectual humility and a constant awareness of how easy it is be misled in the pursuit of knowledge and understanding. So reasonable people are careful and cautious thinkers. Good critical thinkers get well acquainted with lots of logical fallacies, mistakes in reasoning, in order to avoid known pitfalls in thinking. The rational believer doesn't let ego and willfulness get in the way of her evaluation of claims and reasons. The reasonable or rational believer doesn't force things, but rather yields to best reasons and evidence. I won't delve into stock stereotypes here except to point out that being amenable to other perspectives, intellectually humble and cautious, and yielding to the better ideas don't stereotypically characterize maleness, authority or whiteness. But they do characterize reasonableness and rationality.

We started with two metaphysical assumptions here. Namely that we have a shared reality and that we each have limited and fallible experience of that reality. To this we've added a few definitional remarks about truth, rationality, reason, belief, subjectivity and objectivity. And we've reasoned a bit on the basis of these assumptions and definitions. In the definitional remarks I've tried to lay out standard philosophical usage clearly and straightforwardly.

The reason it is good to understand truth, rationality etc. in the manner I've laid out here is that it facilitates clearer communication and understanding of our diverse experiences and diverse ways of thinking. This allows us to cooperatively improve our ways of thinking and our limited grasp of what's true. And as a result of this, we are empowered to act more effectively, avoid hazards in our interactions, appreciate each other more significantly, and enjoy things.

Some will be tempted to object to what I've laid out here on the grounds that people are free to define things like truth and rationality as they please. In a sense, people are free to do so. Nobody has the power to prevent it. All the other concepts we might be tempted to attach to these words are out there and as a linguistic community of a couple

people or of a couple billion we could agree to name those concepts as we please. But to insist on defining things as we like amounts to the privatization of language, with the primary result of undermining our capacity to communicate with one another and understand each other in the limited ways that are open to us. I am no fan of such hyper-individualism. While we could quibble about how to define truth and rationality, the only result of this would be to talk about something else instead. Something other than how we stand as subjects to each other and our shared reality.

Philosophy as Inquiry

Philosophy is a branch of inquiry. Inquiry or investigation is the effort to figure something out, to get at the truth of some matter. Fruitful inquiry that produces lasting knowledge and understanding is typically a community affair. As individuals, all we have to work with is our own individual point of view based on our limited experience and whatever thought processes we are used to or comfortable with. In this predicament, an individual has only new evidence to help check for errors. Worse, many of us are quite adept at only noticing the evidence that supports what we already think, which precisely misses all the evidence that might check for errors in our thinking. This is the problem known as confirmation bias. So, as individuals we are stuck with very limited evidence and little significant error checking. But as a community of inquirers, we can pool our diverse evidence and review our various thought processes. This is how inquiry has always proceeded. The current understanding in an area of science or philosophy is the result of a great many conversations, sometimes going back centuries or millennia, some in person, some in print, some on YouTube.

Once a philosophical position is considered, we want to ask what arguments can be advanced in support of or against that position. In order to get at the best reasons, it is vital that our community of inquirers include people with diverse perspectives and diverse ways of thinking. Otherwise, we are liable to miss important evidence, overlook good reasons, or fail to find flaws in the arguments we do consider. A community of like-minded people will tend to share the same blind spots and hence ignore illuminating alternative perspectives and neglect the error checking these can provide.

We then want to examine the quality of the arguments for and against a given position. Evaluating flawed arguments often points the way towards other arguments and the process of formulating, clarifying, and evaluating arguments continues.

This method of question and answer in which we recursively formulate, clarify, and evaluate arguments is known as **dialectic**. Dialectic looks a lot like debate, but a big difference lies in the respective goals of the two activities. The goal of a debate is to win by persuading an audience that your position is right and your opponent's is wrong. Dialectic, on the other hand, is aimed at inquiry. The goal is to learn something new about the issue under discussion. Unlike debate, in dialectic your sharpest critic is your

best friend. Critical evaluation of your ideas and arguments brings new evidence and reasoning to light. The person you disagree with on a philosophical issue is often the person you stand to learn the most from (and this doesn't necessarily depend on which of you is closer to the truth of the matter).

Dialectic is sometimes referred to as the Socratic Method after the famous originator of this systematic style of inquiry. We will get introduced to some of Plato's dialogues chronicling the exploits of Socrates later. This will give you a good sense for how the Socratic Method works. Then watch for how the Socratic Method is deployed throughout the rest of the course.

The Fruits of Inquiry

We come to know what is true through inquiry. Sometimes this is as straightforward as making some observations. I know that it is sunrise by looking out the window. Sometimes inquiry is an involved process of formulating questions, identifying possible answers, formulating arguments that bear on these and then critically evaluating the arguments in light of whatever evidence we have to work with. The steps in this process may be repeated or elaborated as needed depending on the complexity of the issues raised.

Sometimes inquiry fails to yield definitive knowledge. Sometimes we don't have the evidence we need to settle an issue. And sometimes it is not so clear how to reason well from the evidence we have. As we heard from Bertrand Russell last week, philosophers often fail to uncover the truth of issues they inquire into. So, where inquiry yields no definitive right answer, what's the point of inquiry?

Inquiry bears many fruits even when it doesn't yield final answers. Inquiry can help us:

- clarify our questions
- distinguish different if closely related issues
- identify the plausible answers
- rule out some wrong answers
- appreciate the implications of some possible answers for other related issues
- increase our understanding of issues by doing some or all of the above

As we will see throughout the course, philosophical inquiry often follows a dialectical pattern where we propose a view, offer arguments for that view, critically evaluate the arguments, learn from our mistakes, rinse and repeat. Inquiry proceeds incrementally through a dialectical process of trial and error. As the Muslim philosopher, Alhazan, put it (around 1025):

The seeker after the truth is not one who studies the writings of the ancients and, following his natural disposition, puts his trust in them, but rather the one who suspects his faith in them and questions what he gathers from them, the one who

submits to argument and demonstration, and not to the sayings of a human being whose nature is fraught with all kinds of imperfection and deficiency. Thus, the duty of the man who investigates the writings of scientists, if learning the truth is his goal, is to make himself an enemy of all that he reads, and, applying his mind to the core and margins of its content, attack it from every side. He should also suspect himself as he performs his critical examination of it, so that he may avoid falling into either prejudice or leniency

Notice in this rather militarized analogy that the discovery of truth happens when your attack fails and you “submit to argument and demonstration,” but not to human authority. The real action in this iterated process of dialectical inquiry happens in formulating and evaluating arguments. We’ll get to this shortly, but first I want us to examine the personal traits and social conditions that lead to fruitful reasoning based on our diverse perspectives as ways of thinking.

Conditions for Critical Thinking

We’ve discussed how inquiry draws us towards knowledge and understanding through sharing and critically reviewing the ideas and arguments that emerge from diverse perspectives, experiences and ways of thinking. But this doesn’t happen automatically. Diverse perspectives and ways of thinking can also drive polarization, conflict and hostility. In this section we want to explore how diversity can be harnessed to mutual benefit rather than mutual destruction. We’ve already alluded to a few of these conditions. Here I will list and characterize them more specifically. We’ll start with the personal traits of effective critical thinkers, or just reasonable people. Then we’ll consider the social conditions to be found in communities populated by reasonable people.

- **Fallibilism:** We mentioned earlier that as subjects we are fallible beings. Our evidence is limited and we are liable to make mistakes in reasoning. Given our nature as fallible subjects, we should never be entirely convinced that we have settled a matter once and for all. To be completely convinced we are right would lead us to neglect any further evidence and argument that might warrant some revision of our views.
- **Intellectual Humility:** Closely related to the idea of fallibilism is intellectual humility. Intellectual humility goes beyond merely recognizing our capacity for error. Fallibilism is not directly concerned with our attitudes towards others and their views and thought processes. Intellectual humility does concern these social factors. The intellectually humble person will keep ego out of their engagement with other people in inquiry. Pride and celebration of your excellence is fine in competitive contexts, but inquiry isn’t a competition, it’s a cooperative activity where respect for others is critical. Arrogance and pride are liable to drive others from the project of inquiry with

the result of losing their insights and perspectives. Of course, there are times when someone understands more than others and it may be tempting to see arrogance in expertise. Dismissing expertise as arrogance, however, will be a failure of intellectual humility itself. Seeing arrogance in expertise is a self-protective way of propping up one's ego by judging another as flawed instead of trying to understand them and make good use of a learning opportunity. Bear in mind that genuine expertise is only acquired through the exercise of intellectual humility. This may be hard to see in people who have some hard-earned expertise, but even the smartest among us only move past ignorance by humbly yielding to the better argument.

- **Open Mindedness:** The open-minded person is open to fairly evaluating the reasons and evidence. Note that open mindedness focuses on our openness to reasons and evidence. A popular but misguided conception of open mindedness is that we should never have much confidence in our own beliefs but always grant that we are just as likely to be wrong as someone who disagrees with us. The problem with way of thinking about open mindedness is that the person who knows what they are talking about and holds a view with some conviction as a result of rigorous inquiry would not count as open minded. You might, for instance, encounter a climate change skeptic alleging that climate scientists are not open minded because they are unwilling to consider the possibility the warming of our climate is the result of sun spots. This is a fallacious attempt to undermine the science. We should hold our beliefs with as much conviction as the best available reasons and evidence warrant. Often our reasons justify high levels of confidence, if not absolute confidence (see fallibilism).
- **Intellectual Courage:** Reasonable people, being open minded and intellectually humble, take the risk of discovering that they have things wrong once in a while. This can be hard. It's generally not pleasant to find to find that you are mistaken. It takes intellectual courage to bear this risk with grace. It helps to have a sense of humor here. It's best if curiosity and delight in discovery outweigh the dread we often feel about getting things wrong. But while critical thinking involves a degree of intellectual risk, it should not involve putting your personal safety on the line in any way. Critical thinkers attack ideas and arguments, not each other. If somebody attacks you, they are not being reasonable. It is possible for a person to feel attacked if they self-identify with an idea that comes under scrutiny. But feeling personally attacked when an idea you like faces criticism would be a failure of intellectual humility that results from investing ego into something that isn't you. You are not your ideas. You can change your mind about something without being personally harmed. When reasonable people do change their minds, it will not be because any other person is dominating or compelling them. Reasonable people change their own minds in response to compelling reasons, not domineering people.

- **Perseverance:** Even once we've acquired the traits discussed so far, clarifying and evaluating arguments can be challenging and frustrating work. You might feel this way about some things some things you've already read in this text. For an ounce of encouragement, bear in mind that confusion is often what it feels like to grow intellectually. Of course, sometimes things are confusing because they just don't make sense. But things that do make sense can feel confusing when they are novel, abstract or just complicated. Stick with it. That confusion is what it feels like to grow new neural pathways. You'll be smarter if you see it through. Even after 40 years of studying philosophy, I sometimes find myself feeling lost and confused in my first pass at reading the work of a philosopher I haven't studied before. Then in the second reading things will begin to make sense. Take some rest between passes. It also helps enormously to take notes on how terms are defined and how arguments are structured. Your brain will continue to sort things out even when you aren't actively reflecting on the material. By the third or fourth pass, maybe over the course of a few days, rich and clear understanding will emerge and you'll be wiser than you were before.

Now let's consider what a community of reasonable people who uphold these intellectual virtues will look like. I think it will be a community characterized by freedom from domination, tolerance and respect for diverse others, good humored civility, a healthy political capacity to deal with shared problems and challenges cooperatively and effectively, and intimacy in friendship. Let's consider each of these in more detail.

- **Freedom from domination:** Critical thinking provides a way of exploring, understanding and sometimes resolving differences between people. This is an alternative to bullying, manipulation, deceit and domination. Again, critical thinkers are responsive to good reasoning and they cultivate intellectual defenses against rhetorical bullying and propaganda aimed at social control. Critical thinkers will resist dominating attempts to by-pass their own intellectual capacities through manipulation or deceit.
- **Tolerance and Respect for Diverse Others:** People who recognize their own fallibility and value intellectual humility will recognize that intolerance bars others from sharing their evidence and argument. This will introduce blind spots in inquiry and frustrate attempts to understand things and figure things out. Likewise disrespectful treatment of others is liable to drive them from participating in inquiry with the same result of ignoring potentially important evidence and argument. Intolerance and disrespectful treatment of others is literally a recipe for ignorance.
- **Politics:** I'm sure you have noticed how divisive politics is in America at the moment. Passionate conflict in politics often reflects a struggle for power aimed at sustaining or overcoming oppressive domination. But even here, conflict is driven and amplified by poor critical thinking. The political polarization in we current see in

America is the result of people refusing to try to understand each other and evaluate each other's reasons and perspectives fairly. I'm afraid a great many Americans have become unreasonable people, disastrously poor critical thinkers. If we were better able to understand and evaluate each other's perspectives, we would be much more capable of finding common ground in addressing our shared problems. If we were better able to identify fallacies, mistakes in reasoning, we would be much less vulnerable to manipulation that divides us and undermines mutual understanding.

- **Friendship:** There may be no more basic human need than the need to be loved. As subjects, we are doomed to a sort of isolation. No other person, not matter how well they know you and care for you can share your subjectivity. We can only hope to understand each other to limited degrees. But I'd submit that the drive to charitably understand another people is itself a form of love.

This probably sounds idealistic to the point of being unrealistic. That is understandable given the current state of our world. We face multiple crises from political dysfunction to climate change and this engenders a great deal of fear and anxiety. In this state, critical thinking is not just intellectually challenging, but it is likely to feel emotionally remote as well. When people are fearful and anxious it natural to seek security in the familiar and defend that against all intrusions. The need for intellectual courage is all the more dire and may seem to carry with it a need for emotional courage just when this seems least available. What I want to suggest here, is that we can seek comfort and security not only in the familiar, but also in the project of building communities of critical thinkers. This obviously starts with cultivating our own critical thinking skills. And this may require loosening our grip on ideological security blankets. But as we saw last week in connection with Russell, clinging to opinions as a security blanket doesn't really provide security. A better strategy is to seek comfort and security in friends and loved ones. Critical thinking provides an avenue to expanding your community of friends and loved ones even across great differences of perspective.

2. Critical Thinking II: Logic

Philosophers, as we've previously mentioned, are mainly in the business of formulating, clarifying and evaluating arguments. This is how inquiry proceeds. In any realm of inquiry, this is how we determine what is true, when we can. An argument is a reason for thinking something is true. An argument consists of a set of premises which work together to provide a reason for accepting a conclusion as true. In this chapter we will get introduced to the basic standards and procedures for formulating, clarifying and evaluating arguments.

We've introduced the idea of an argument as a reason for believing something and most the chapter will focus on this primary function of arguments. But arguments are multifunction tools in inquiry and we will also want to discuss their various other uses along the way. Here are a few:

Arguments can be useful for

- Providing a reason for thinking their conclusions are true
- Clarifying our reasons
- Teasing out false premises
- Clarifying our own conceptual understanding
- Recognizing gaps on our own reasoning
- Understanding the views of others

We'll have to say more about how to clarify and evaluate arguments before explaining these points.

Arguments

The way to determine whether a claim is true or false, when this is possible, is to evaluate the evidence and argument for and against it. Sometimes good reasons take the form of simple observations. I have a good reason for thinking my bicycle has a flat tire when I see the tire sagging on the rim. But often the business of identifying and evaluating reasons is a bit more involved.

An argument is a reason for taking something to be true. Arguments consist of two or more claims, one of which is a conclusion. The conclusion is the claim the argument aims to establish as true. The other claims, there can be one or many, are the premises. The premises of an argument taken together are offered as a reason for believing its conclusion to be true.

Some arguments provide better reasons for believing their conclusions than others. In case you have any doubt about that, consider the following examples:

1. Sam is a line cook.
2. Line cooks generally have good of kitchen skills.
3. So, Sam can probably cook well.

1. Sam is a line cook.
2. Line cooks usually aren't paid very well.
3. So, Sam is probably a millionaire.

The premises in the first argument provide pretty good support for thinking Sam can cook well. That is, assuming the premises in the first argument are true, we have a good reason to think that its conclusion is true. The premises in the second argument constitute a pretty poor reason to think Sam is a millionaire. So, whether or not the premises of an argument support its conclusion is one important factor in evaluating an argument.

Now consider these examples:

1. Boston is in Massachusetts.
2. Massachusetts is east of the Rockies.
3. So, Boston is east of the Rockies.

1. Boston is in California.
2. California is west of the Rockies.
3. So, Boston is west of the Rockies.

Again, the first of these two arguments looks good, the second not so much. But the problem with the second of these arguments is different. The premises of both arguments provide good support for the conclusion. That is, in both arguments, if the premises were true, we'd have good reason for accepting the conclusion. In fact, for both arguments, if the premises were true, the conclusion would have to be true. So, in both of these arguments we have a good relation of logical support between the premises and the conclusion. But the first premise of the second argument just isn't true. Boston is not in California. So, the latter pair of arguments suggests another key issue for evaluating arguments. Good arguments have true premises.

That is pretty much it. A good argument is an argument that has true premises that support its conclusion. So, evaluating an argument involves these two steps:

- Determine whether or not the premises are true.
- Determine whether or not the premises support the conclusion (that is, whether we have grounds to think the conclusion is true if all of the premises are true).

Determining whether an argument's premises are true may involve evaluating further arguments in support of those premises. An argument might be the last link in a long chain of reasoning. In this case, the quality of the argument depends on the whole chain. And since arguments can have multiple premises, each of which might be supported by further arguments, evaluating an argument might be more involved yet, since its conclusion is really supported by a rich network of reasoning, not just one link and then another. While the potential for complication should be clear, the basic idea should be pretty familiar. Think of the regress of "why" questions many of us tormented our parents with as children. Even at a young age we understood that the reasons for believing one thing can depend on the reasons for believing a great many other things.

However involved the network of reasons supporting a given conclusion might be, it seems that there must be some starting points. That is, it seems there must be some reasons for believing things that don't themselves need to be justified in terms of further reasons. Otherwise, the network of supporting reasons would go on without end. The issue we are facing here is one of identifying the ultimate foundations of knowledge and justified belief. This is a big epistemological issue and we will return to it later in the course. For now, let's consider one potential answer we are already familiar with. In the sciences our complex chains of reasoning seem to proceed from the evidence of the senses. We think that evidence provides the foundation for our edifice of scientific knowledge. Sounds great for science, but where does this leave philosophy? Does philosophy entirely lack evidence on which its reasoning can be based?

Philosophy does have a kind of evidence to work from and that evidence is provided by philosophical problems. When we encounter a problem in philosophy this often tells us that the principles and assumptions that generate that problem can't all be correct. This might seem like just a subtle clue that leaves us far from solving the big mysteries. But clues are evidence just the same. As we will discuss in our chapter on the philosophy of science, science doesn't really have it much easier. Sensory evidence by itself doesn't tell us as much about the nature of the world as we often suppose. Scientific evidence provides clues, but there remains a good deal of problem solving to do in science as well as in philosophy.

So, we can assess the truth or falsity of the premises of an argument by examining evidence or by evaluating further argument in support of the premises. Now we will turn to the other step in evaluating arguments and consider the ways in which premises can support or fail to support their conclusions. The question of support is distinct from the question of whether the premises are true. The reason one of our arguments about Sam the line cook was good but not the other had nothing to do with false premises. We can grant that the premises in both arguments were true. The difference had to do with whether the premises provided good support of the conclusion. When we ask whether some premises support a conclusion, we are asking whether we would have good

grounds for accepting the conclusion if we assume that the premises are true. It is important that we keep the two steps in evaluating arguments distinct in our minds. When we evaluate arguments holistically, as people often do, we wind up accepting or rejecting arguments based on how we feel about them overall without looking into whether the premises of the arguments really support the conclusions we draw. This is one of the ways we fall victim to confirmation bias, by endorsing just the arguments that point towards the conclusions we like without scrutinizing the logic of the argument.

Consider again the two good arguments in our examples above:

1. Sam is a line cook.
2. Line cooks generally have good of kitchen skills.
3. So, Sam can probably cook well.

In this example the premises do support the conclusion. We have pretty good reason to think Sam can cook well if he is a line cook. But these premises don't guarantee that Sam can cook well. It might be his first day on the job. He might be a really lousy line cook. Or he might be a breakfast cook and pretty useless in the kitchen beyond frying eggs and making hash browns. Still, the premises of this argument would give us good reason for trusting him with dinner. The premises being true would make it pretty likely he'd feed us well.

Now consider this one again:

1. Boston is in Massachusetts.
2. Massachusetts is east of the Rockies.
3. So, Boston is east of the Rockies.

In this argument the premises don't just make the conclusion likely. The premises being true would guarantee the truth of the conclusion. These two examples point us towards our two standards of support, deductive validity and inductive strength. A deductively valid argument is one where the premises, if they are true, would guarantee the truth conclusion. The support relation in the case of deductively valid arguments is logically necessary. Inductively strong arguments are arguments where the premises, if they are true, would provide good reasons for thinking the conclusion is true. But good reasons in inductively strong arguments are a matter of probability, not necessity. A strong inductive argument with true premises doesn't guarantee the truth of the conclusion.

Deductive Validity

The deductive standard of support is validity. An argument counts as deductive whenever its aiming at validity. Deductive validity is the strictest standard of support we can uphold. In a deductively valid argument, the truth of the premises guarantees the truth of the conclusion. Here are two equivalent definitions of deductive validity:

(D) A valid argument is an argument where if its premises are true, then its conclusion must be true.

(D') A valid argument is an argument where it is not possible for all of its premises to be true and its conclusion false.

Here are a few examples of deductively valid arguments

1. If Socrates is human, then Socrates is mortal
2. Socrates is a human.
3. Therefore, Socrates is mortal

1. All monkeys are primates
2. All primates are mammals
3. So, all monkeys are mammals

If you think about these two examples for a moment, it should be clear that there is no possible way for the premises to all be true and the conclusion false. The truth of the conclusion is guaranteed by the truth of the premises. In contrast, the following argument is not valid:

1. If Sue misses her plane, she will be late for the conference.
2. Sue is late for the conference.
3. Therefore, she missed her plane.

Again, to say that an argument is deductively valid is to say that it is impossible for all of its premises to be true and its conclusion to be false. To see why the last argument is not valid, try to think of a possible scenario that makes both of the premises true and the conclusion false. One scenario is where Sue catches her plane, but her cab from the airport gets stuck in traffic. Another would be where Sue makes her flight, but the plane is delayed due to bad weather. If we can think of any possible way for the premises of an argument to be true and its conclusion false, then we have shown that the conclusion does not deductively follow from the premises. That is, we've shown that the argument is not valid.

Our intuitive test for validity is to think about whether it is possible for the argument's premises to be true and its conclusion to be false. A key point to notice here is that validity is not directly about the truth or falsity of the premises or the conclusion. The concept of validity is really a concept about what is and isn't logically possible. A deductively valid argument may or may not have true premises. Consider this argument:

1. All planets are stars.
2. All stars are bodies that shine steadily.

3. All planets are bodies that shine steadily.

Both of the premises in this argument are false, but the argument is still valid. Suppose, contrary to fact, that the premises were true. The conclusion would have to be true if this were the case. Validity isn't about whether the premises or the conclusion are in fact true. It is only about whether the conclusion logically follows from the premises.

Given this, a deductively valid argument only provides one with a good reason for believing its conclusion *if its premises are true*. If a deductively valid argument has all true premises, we say that it is deductively sound. For an argument to be deductively sound is one way for it to pass both steps for evaluating arguments. A deductively sound argument has all true premises that support its conclusion.

The deductive arguments we've looked at here are pretty intuitive. We only need to think about whether the conclusion could be false even if the premises were true. But most deductive arguments are not so obvious. Logic is the science of deductive validity. Philosophy has made some historic advances in logic over the past century. Bertrand Russell, who we got acquainted with in the last chapter, was among the key contributors to developments in logic early in the 20th century.

Logical Form

Since Aristotle, the first major logician, it's been recognized that deductive validity is a matter of an argument's logical form. We can display an argument's logical form by replacing all but the logically operative vocabulary with symbols (we'll use capital letters for this). So, consider the logical form of a few of our examples so far.

1. All planets are stars.
2. All stars are bodies that shine steadily.
3. All planets are bodies that shine steadily.

This argument has the following form:

1. All P are S
2. All S are B
3. All P are B

Any argument that has this logical form will be valid. So,

1. All fish are vertebrates.
2. All vertebrates are animals.
3. So, all fish are animals.

Remember, validity is just a standard of support. Validity does not assume true premises or a true conclusion. So even though it sounds a bit “off,” this argument is also valid:

1. All red things are bricks,
2. All bricks are rocket ships.
3. So, all red things are rocket ships.

Of course, this argument sounds silly. Both premises are ridiculously false. But then any possible world where both premises are true would be a possible world where all red things are rocket ships. The argument is valid in virtue of its valid logical form. Now consider this familiar argument:

1. If Socrates is human, then Socrates is mortal
2. Socrates is a human.
3. Therefore, Socrates is mortal

This argument has the following logical form:

1. If H, then M
2. H
3. M

Similarly, any argument that has this logical form will be valid. Plug any declarative sentences you like in for H and M and you will have a valid argument. The premises might be false, or even absurd, but it will remain the case that any way the world could be that makes both premises true will also make the conclusion true. Once you appreciate how deductive validity is a function of the logical form of an argument, it soon becomes clear that a valid argument can be constructed for any possible conclusion, true, false, or completely absurd. So, for instance:

1. If pigs fly, then the oceans will dry up.
2. Pigs fly
3. Therefore, the oceans will dry up.

So, you might be wondering what the point of all this silliness is. It's partly to limber up your logical sense and help you recognize that logical validity is only about what follows from what, not about what is in fact true or false. Of course, the oceans aren't going to dry up. But if both premises were true, then the conclusion would follow logically and also be true. But there is a further point to the hypothetical silliness. The fact that the conclusion of the “pigs fly” argument is absurdly false is a good indicator that at least some of the premises of this valid argument are also false. And this is a very useful

thing to recognize. To see this, let's look at another valid argument pattern that captures what we've just said about the pigs fly argument:

1. If P, then C
2. Not C
3. So, not P

This is a valid pattern of reasoning that we use routinely. For instance:

1. If I have milk, then it will be in the fridge
2. There's no milk in the fridge
3. So, I am out of milk.

Now notice how we used this pattern of reasoning in our analysis of the "pigs fly" argument. It is valid, which means that if its premises are all true, then its conclusion is true. But obviously, its conclusion isn't true. So, its premises are not all true.

Reductio ad Absurdum

I mentioned near the beginning of this chapter that arguments are multifunctional tools in inquiry. Arguments aren't always used directly to show the truth of some conclusion. As we've just seen, the concept of a valid argument can be used to tease out falsity in the premises. For instance, we might consider a claim that sounds pretty good and ask what follows from that claim deductively. What conclusion could we validly argue for on the basis of that claim? If we find that by deductively valid reasoning we can get from our claim that sounds pretty good to an absurd conclusion, then we have shown that our starting point, the claim that sounded pretty good, is false. This strategy is known as *reductio ad absurdum*, which is a handy bit of Latin for "reducing to absurdity." We can use this strategy to test an idea for problems by considering what follows from that idea by valid argument and making sure it doesn't lead to anything obviously false or absurd.

To illustrate *reductio ad absurdum*, let's consider a view we mentioned briefly in the first chapter of this text about the nature of morality. A view that many people find attractive: moral relativism. According to moral relativism, there are no objective moral standards, rather morality is relative to groups depending on what is considered right in that group. When we consider what follows from moral relativism deductively, we wind up with some pretty unsavory results. The first premise in this argument is just a statement of moral relativism as a view about the nature of morality. From here, bad things start to happen.

1. If a society considers something morally good, then it is morally good (relative to that society).
2. Nazi Germany considered the extermination of Jewish people good.
3. The extermination of Jewish people was good (relative to Nazi Germany).

The argument here is valid. It's logical form is a minor variation on a valid pattern we examined above. If the premises are true, then the conclusion must be true. This means that if the conclusion is false, at least one of the premises must be false. Note that the conclusion here is not about what the Nazis considered to be good, it's about what *is* good in the only sense that matters according to moral relativism. Since the conclusion of this argument is obviously false, not to mention horrible, and the second premise is a matter of historical fact, moral relativism must be false. Here we have reasoned validly from a view about the nature of morality that many people find attractive to a conclusion that is obviously absurdly false and horrible. A view about the nature of morality that has obviously and horribly bad logical consequences has got to be false. Moral relativism reduces to absurdity, *reductio ad absurdum*. We will get to examine moral relativism in greater detail when we get around to studying ethics. The point here is just to see how deductive argument is useful not just for getting at true conclusions, but also for teasing out false premises using the strategy we'd just identified as *reductio ad absurdum*.

Revealing Hidden Assumptions

Another very helpful function of valid argumentation is as an aid to revealing hidden assumptions. In everyday life we don't generally formulate formally valid arguments when we give reasons for what we believe. And often these unmentioned premises are where our biases hide. A good understanding of deductive validity can help us bring these hidden assumptions and biases to light. For example:

1. Every story I hear about politician X on Facebook says he is doing a terrible job.
2. So, Politician X is doing a terrible job.

This is not a valid argument as it stands. We would have a valid argument if we added a premise as follows.

1. Every story I hear about politician X on Facebook says he's doing a terrible job
2. If every story I hear about politician X on Facebook says he's doing a terrible job, then politician X is doing a terrible job.
3. So, Politician X is doing a terrible job.

Now we have a valid argument, but adding the premise required to have a valid argument reveals a hidden assumption that, as some of you probably know, we have reason to doubt. Facebook uses AI and algorithms to customize what you see in order to maximize engagement. It turns out that anger is very engaging. As a result, Facebook tends to feed you news stories that will stoke your anger. If you have a history of clicking and commenting on stories that say awful things about politician X or others of his political persuasion, Facebook will load your news feed with more articles that say awful things about politician X. The same goes for Google, YouTube and most search engines and social media platforms. The reason Facebook keeps showing you

stories about what a terrible job politician X is doing isn't that politician X is actually doing a terrible job. Rather it's that the AI behind Facebook algorithms knows that stories like this will keep you glued to the platform, posting angry comments about politician X.

So, formulating deductively valid arguments brings our assumptions to light where they can be scrutinized for truth or reasonableness. A good understanding of deductive validity can be very useful in identifying and addressing our usually unspoken assumptions and biases (well, perhaps in the example we just considered, the bias lies mainly in the AI driven algorithms employed by Facebook).

Clarity

For reasons we just been discussing, a good understanding of deductive validity can help you clarify your own reasons and express them clearly to others. If you can recognize when an argument you find appealing has some deductive gaps in it, your understanding of validity will guide you in filling in those gaps. Assuming the argument is a good one, you will then have a clearer understanding of it and be able to express your reason more clearly to others. Of course, as just discussed, if your argument is not so great, your understanding of validity can alert you to this and perhaps guide you in formulating a better argument. All of this applies to understanding the arguments offered by others. When others formulate their reasons in incomplete, less than valid ways, your understanding of validity can guide you in identifying the questions you'd want to ask in reconstructing a more complete version of your friend's argument.

Charitable Interpretation

A good understanding of deductive validity will help you formulate the clearest and best possible arguments for your view. It will also help you formulate and appreciate the best possible arguments for views you oppose. Formulating the best possible interpretation of an argument for opposing views is what we call "charitable interpretation. In the dialectical spirit of cooperatively working towards truth and reasonableness, it is best to be charitable in filling out your reconstruction of another's reasons. While helping to clarify arguments is a kindness, this isn't really the point of charitable interpretation. Finding faults with bad arguments for a view you disagree with doesn't really undermine that view, it just undermines bad arguments for that view. If you have good reasons for rejecting a view, you should aim to identify the flaws in the best possible arguments for the view you think is wrong. Trying to make the opposing view sound outrageous or ridiculous will only lead you into straw man attacks (see discussion of the fallacy below). The strongest argument you can offer against a view you oppose is not the argument that makes the view sound outrageous, but the argument that targets the best interpretation of the opposing view and the best possible arguments for it. Being a reasonable and effective critical thinker calls for charitable interpretation of opposing

views and the arguments for them, not just out a sense of fair mindedness, good as that may be, but also in order to be the most effective critic of the view you oppose.

In the last few sections, I've tried to characterize a few useful functions for deductive argumentation beyond merely trying to give reasons for thinking that something is true. While sound arguments, arguments that are both valid and have all true premises, provide good reasons for accepting their conclusions as true, valid argument isn't only useful for directly getting at the truth. A good understanding of validity is useful for clarifying reasoning and bringing hidden assumptions to the surface. It can be useful for drawing our attention to false premises. And it can help us make our criticism of views we oppose more effective by targeting the best versions of those views. The path to knowing truths and understanding issues is usually not a straight paved sidewalk. It takes some skill to recognize the switchbacks and stay on the trail. A good understanding of deductive validity is powerful guide.

Inductive Strength

I won't have as much to say about inductive strength and cogency since you will already be more familiar with it from science classes and because philosophy trades more in deductive argument. Let's start with our example argument from above:

1. Sam is a line cook.
2. Line cooks generally have good of kitchen skills.
3. So, Sam can probably cook well.

This is a decent argument. The premises do support the conclusion. And yet it might be that both premises are true and the conclusion is false. Sam could be a brand new cook, hired because he's the manager's son, but also someone who has never cooked in his life. Many arguments give us good reasons for accepting their conclusions even when true premises don't guarantee the truth of the conclusion. This suggests that we need another standard of support for arguments that aim at giving us pretty good but not absolutely compelling grounds for accepting their conclusions. And this standard of support is called inductive strength. Here are two equivalent ways of defining inductive strength:

- (I) An inductively strong argument is an argument in which if its premises are true, its conclusion is probably to be true.
- (I') An inductively strong argument is an argument in which it is improbable that its conclusion is false given that its premises are true.

If you look again at the earlier definitions for deductive validity you will find a good deal of similarity. The only difference is in the use of the words "probably" rather than "must

be" in the first definition, and "improbable" rather than "impossible" in the second. This is a big difference. As in the case of validity, when we say that an argument is strong, we are not assuming that its premises are true. We are only claiming that *if* the premises are true then the conclusion is likely to be true. Corresponding to the notion of deductive soundness, an inductive argument that is both strong and has true premises is called a *cogent* inductive argument. Unlike the case if deductively sound arguments, it is possible for an inductively cogent argument to have true premises and a false conclusion.

What makes an argument an inductive argument is that it is aiming at the standard of inductive strength. Similarly, what makes an argument a deductive argument is aiming at validity. Students frequently ask if an invalid deductive argument can be considered inductively strong. Generally, not. The targets are different. Missing the target of deductive validity doesn't make an argument inductively strong. Invalid deductive arguments are generally just bad arguments. Once in a while there will be a decent inductive argument that looks a bit similar, but not typically. Deductive and inductive refer to different kinds of reasoning.

Lots of good reasons for holding a belief fall short of the standard of deductive validity. The sort of reasoning you were taught as "the scientific method" is inductive reasoning. As it is taught in high school, the scientific method consists of formulating a general hypothesis and testing it against a large sampling of data. If the data is consistent with the hypothesis, then the hypothesis is considered confirmed by the data. Here a limited amount of evidence is taken to support a broader more general hypothesis. In the simplest case, inductive reasoning involves inferring that something is generally the case from a pattern observed in a limited number of cases. For instance, if we were to conduct a poll of 1000 Seattle voters and 600 of them claimed to be Democrats, then we could inductively infer that 60% of the voters in Seattle are Democrats. The results of the poll give a pretty good reason to think that around 60% of the voters in Seattle are Democrats. But the results of the poll don't guarantee this conclusion. It is possible that only 50% of the voters in Seattle are Democrats and Democrats were, just by luck, over represented in the 1000 cases we considered.

When evaluating deductive arguments for validity we ask if it is possible for the premises to be true and the conclusion to be false. This is either possible or it isn't. Possibility does not admit of degrees. But probability does. The truth of the conclusion of an inductive argument can be probable to a greater or lesser degree. An argument either is or isn't valid. But inductive arguments can be more or less strong. We can identify a few factors that bear on the degree of strength an inductive argument has. One is how much evidence we have looked at before inductively generalizing. Our inductive argument above would be stronger if we drew our conclusion from a poll of 100,000 Seattle voters, for instance. And it would be much weaker if we had only polled

100. Also, the strength of an inductive argument depends on the degree to which the observed cases represent the makeup of the broader class of cases. So our inductive argument will be stronger if we randomly select our 1000 voters from the Seattle phone book than if they are selected from the Ballard phone book (Ballard being a notably liberal neighborhood within Seattle).

So far, we've only discussed inductive generalization, where we identify a pattern in a limited number of cases and draw a more general conclusion about a broader class of cases. Inductive argument comes in other varieties as well. In the example we started with about Sam the line cook, we inductively inferred a prediction about Sam based on a known pattern in a broader class of cases. Argument from analogy is another variety of inductive reasoning that can be quite strong. For instance, I know that my housecat is very similar to cougars in the wild. Knowing that my cat can jump great heights, it would be reasonable to expect that by analogy, or based on this similarity, cougars can jump well too.

There are further varieties of argument that aim at the standard of inductive strength, but we will discuss just one more in detail now. Abduction is inference to the best explanation. Detective work provides a good example of abductive argument. When Holmes discovers Moriarty's favorite brand of cigar and a bullet of the sort fired by Moriarty's gun at a murder scene, inference to the best explanation suggests that Moriarty was the killer. That Moriarty committed the murder provides the overall best explanation of the various facts of the case.

The 19th century American pragmatist and logician, Charles Sanders Peirce offers the *Surprise Principle* as a method for evaluating abductive arguments. According to the surprise principle, we should count one explanation as better than competing explanations if it would render the facts we are trying to explain less surprising than competing explanations. The various clues in the murder case are among the facts we want explained. The presence of the cigar and the bullet casing at the murder scene is much less surprising if Moriarty committed the murder than if the maid did it. Inference to the best explanation aims at strength. So a strong abductive argument in this case needn't rule out the possibility that the murder was committed by Moriarty's evil twin who convincingly frames his brother. There might an argument against the death penalty lurking nearby. Inference to the best explanation is worth more attention than it often receives. This kind of reasoning is pervasive in philosophy and science, but seldom gets much notice as an integral part of the methods of rational inquiry.

Fallacies

A fallacy is just a mistake in reasoning. Humans are not nearly as rational as we'd like to suppose. In fact, we are so prone to certain sorts of mistakes in reasoning that

philosophers and logicians refer those mistakes by name. In this section we'll get acquainted with a choice selection of all-to-human irrational tendencies. For thorough catalogue of logical fallacies, I'll refer to you [The Fallacy Files](http://www.fallacyfiles.org/taxonomy.html) (<http://www.fallacyfiles.org/taxonomy.html>).

- **Ad hominem:** This fallacy is known is Latin for “against the person.” As the name suggests, *ad hominem* consists of attacking the proponent of a position rather than critically evaluating the reasons offered for the proponent’s position. The reason *ad hominem* is a fallacy is that the attack on an individual is simply not relevant to the quality of the reasoning offered by that person. Attacking the person who offers an argument has nothing to do whether or not the premises of the argument are true or whether they support the conclusion. *Ad hominem* is a particularly rampant and destructive fallacy in our society. What makes it so destructive is that it turns the cooperative social project of inquiry through conversation into polarized verbal combat. This fallacy makes reasonable dialogue impossible while it diverts attention from interesting issues that often could be fruitfully investigated.

Here’s an example of *ad hominem*: A car salesman argues for the quality of an automobile and the potential buyer discounts the argument with the thought that the person is just trying to earn a commission. There may be good reason to think the salesman is just trying to earn a commission. But even if there is, this is irrelevant to the evaluation of the reasons the salesman is offering. The reasons should be evaluated on their own merits. We can imagine a situation where the salesman is just trying to earn a commission and yet he is also making good arguments. Consider a salesman who is not too fond of people and cares little for them except that they earn a commission for him. Otherwise, he is scrupulously honest and a person of moral integrity. In order to reconcile himself with the duties of a sales job, he carefully researches his product and only accepts a sales position with the business that sells the very best. He then sincerely delivers good arguments for the quality of his product, makes lots of money, and dresses well. Perhaps this is an unusual situation. Perhaps this salesman was a philosophy major. The customer who rejects his reasons for buying the car he sells on the *ad hominem* grounds that he is just trying to earn a commission misses an opportunity to buy the best. The moral of the story is just that the salesperson’s motive is logically independent of the quality of his argument. Of course, to know this, you’d want to also read some reviews, study the car yourself and make sure his premises check out.

- **False Dichotomy:** A dichotomy is an either/or choice where this is no third or fourth option. We’ve seen an example of a dichotomy in the contrast between the claim that there is intelligent life on other planets and the claim that there is no intelligent life on other planets. If one option is false then the other is true. There is no third or

fourth possibility. On the other hand, when you go to a restaurant and you are trying to decide between the Impossible Burger or the Caesar Salad, you are probably not facing a dichotomy. You also have the option of having the salmon, or perhaps the fajita. The fallacy of false dichotomy is committed when we are presented with just two options as if these were the only possibilities when in fact there may be a third, fourth or more other possibilities.

So here is a famous example of the false dichotomy fallacy. Shortly after 911, George W. Bush proclaimed, "You are either with us or you are with the terrorists." Not long after that Bush launched a war against Iraq in the name of fighting terrorism. Some Americans protested the invasion of Iraq, arguing that we did not have good reason to feel threatened by that country and that given this, an unjust war would inspire more terrorism than it prevented. People who protested the invasion of Iraq were roundly vilified as "terrorist sympathizers." In fact, critics of the war in Iraq were as opposed to terrorism as the rest of America, they simply doubted that the Bush administration's invasion of Iraq would be an effective way to combat terrorism. As it turns out, Iraq did not have the weapons of mass destruction it was alleged to have and it was not at the time a sponsor of terrorism (though many members of Iraq's disbanded military went on to join terrorist organizations after we invaded Iraq). The long war in Iraq took an enormous toll on human life and well-being both to US and international service personnel but mostly to Iraqi citizens. With a long-term price tag of around 2 trillion for US tax payers it was enormously expensive. The war badly damaged US standing on the international stage. And it inspired a great deal of terrorism, mostly in Iraq and then Syria, often targeting US military personnel but also civilians. Critics of the war here in the US were not either with the Bush administration or with the terrorists. They shared the Bush administration's goal of ending terrorism but opposed its strategy for achieving this goal.

- **Straw Man:** When soldiers fought with bayonets on their rifles, they would train by attacking straw men. Straw men are fairly easy to stab with a bayonet since they don't run away or fight back. But then stabbing a straw man is no victory over an actual opponent. The fallacy of straw man is committed when someone criticizes an easy to attack distortion of an argument or idea rather than the actual view. Like many fallacies, this one can be committed deliberately or inadvertently. In our highly polarized social media environment, it is not uncommon for a disingenuous manipulator to deliberately broadcast a straw man attack (or some other fallacy) all the while knowing that his audience, lacking well developed critical thinking skills, will fall for the manipulation and go on to propagate the bad argument sincerely. This is often how propaganda works.

You may have heard a commonly propagated straw man attack committed against efforts to address climate change. Critics will often charge that people concerned about the environment are really just socialists looking to take our freedom away. There's a lot going on here and its worth pointing out the fallacies are gregarious. It is quite possible to commit more than one at a time. So, you might also notice an element of *ad hominem* in this example where reasons for taking climate seriously get ignored in favor of attacking the people trying to do so. This sheds some light on the old quip that lies travel half way around the world before the truth gets its shoes on. It takes lots more work to diagnose and filter out fallacies than it does to commit and propagate them. But aside from the *ad hominem* attack, this notion that people who want to see action on climate change are just big government lovers includes a straw man distortion of the climate advocates actual views. Yes, addressing climate change is going to require strong government action and regulation of the use of fossil fuels, eventually replacing them altogether. The straw man does look something like the government loving enemy to people who commit this fallacy. The fallacy wouldn't work without this resemblance. But climate advocates don't just love big government and freedom infringing regulation. They rather recognize a dire and pressing systemic global problem and they recognize that it can only be addressed by coordinated government action. Many of them accept the necessity of regulation with reluctance. Private and individual efforts to address climate change are certainly helpful, perhaps also necessary. But the climate problem is not really going to get addressed without effective policy that can only get enacted and administered through government.

There are many more fallacies worth getting familiar with. I'll leave you to explore these on your own. The [The Fallacy Files](#) is a good place to start. I'll wrap up here with a brief mention of confirmation bias. Confirmation bias isn't really a fallacy because it isn't a specific kind of mistake in reasoning. Confirmation bias is the intellectual bad habit of endorsing just the evidence and argument that seems to support the view you already hold. Just about any fallacy can be involved in confirmation bias. Confirmation bias is about what you'd expect to see among people who lack strong critical thinking skills. People who don't know how to evaluate arguments have little else to go on except to prefer the arguments with conclusions they like. The only defense anyone has against confirmation bias is building some critical thinking skills, learning some logic and learning to identify fallacies.