4 | Drawing on Logic

In the last chapter, you were introduced the subject of logic in the context of Aristotle's three modes of persuasion: logos, pathos, and ethos. I hope that discussion clued you into the foundational importance of logical thinking, whether you are choosing who to vote for, deciding whether or not to buy a house or change jobs, or coding a new app.

The subject of logic (or really logics, since there exist today multiple logical systems) is vast, intricate, and, for some, off-putting, especially when it seems to intersect with mathematics. While not dismissing the importance of all these systems (including the mathematical ones), this chapter will draw upon some practical verbal and visual concepts and tools from the vast grab bag of logic that you can put to immediate use to help you debunk bunk, communicate persuasively, and make good choices in life.

So let's start with one type of logic—deductive—and one of the first techniques used to systematize logic: Aristotle's syllogisms.

You've already been introduced to deductive logic and syllogisms in the last chapter. Just to recap, a deductive argument is self-contained in that everything needed to evaluate whether the argument works or not is contained within the argument itself. (We'll get into inductive arguments, which are not so self-sealed, later in this chapter.)

When you write an argument in the form of a syllogism, it must be structured in a very specific way using very specific language. For example, here is a simple syllogism you were introduced to previously that we'll be using to illustrate various concepts of logic described throughout this chapter:

All dogs are animals.
Francine is a dog.
Therefore, Francine is an animal.

In this example, the syllogistic argument contains two premises ("all dogs are animals" and "Francine is a dog") followed by a conclusion ("Francine is an animal", the word "therefore" being used to signal that a conclusion follows). So adding labels to our argument, we get this:

Premise 1: All dogs are animals.
Premise 2: Francine is a dog.
Conclusion: Therefore, Francine is an animal.

So far so good.

Now just to get a little more terminology out of the way, the first premise in this argument is called the *major premise* (it's usually the one that makes the more sweeping statement) and the second one is (you guessed it) the *minor premise*. Because "animals" appears in just the major premise, it is called the *major term*. Similarly, "Francine" is the *minor term* and "dogs" (which appears in both premises) is called the *middle term*.

The statements in a syllogism, including the premises and the conclusion, must all be written in a certain way so that they fit into one of these four categories:

- All A's are B's (called an A statement).
- No A's are B's (called an E statement).
- Some A's are B's (called an I statement).
- Some A's are not B's (called an O statement).

With those preliminaries out of the way, most of us are less concerned with how to label parts of an argument than we are with determining whether the argument is any good. And one of the tests to see if an argument works or not is called *validity*.

So what is validity? Well at the most basic level, an argument is valid if accepting the premises means you must accept the conclusion. So if I ask you to accept the premises of my argument and you agree, and there is no way for you to reject the conclusion (by, for instance, finding a counter-example or otherwise determining the conclusion to be false), then the argument is considered a valid one.

If you look over our animals/dogs/Francine example, it certainly seems as though this argument is valid in that it's pretty hard (OK, it's impossible) to accept the two premises (that all dogs are animals and that Francine is a dog) but then declare it false that Francine must be an animal. But given that not all arguments are this straightforward, we need a system for determining validity that does not rely solely on our intuition.

Aristotle came up with one such system that took advantage of the highly structured nature of the syllogism and determined that if a syllogism were constructed correctly (i.e., if the major, minor, and middle terms were distributed appropriately throughout the premises and conclusion), then certain combinations of A, E, I, and O statements were always valid while others were invalid (or only valid under special circumstances). For example, a correctly structured syllogism consisting just of A statements (called an

AAA syllogism)—like our animals/dogs/Francine example—is valid, while other combinations (like EEE) are not.

For centuries, students used mnemonic devices to remember which of the more than two hundred possible combinations of A, E, I, and O statements were valid arguments and which were invalid. For example, "B**a**rb**a**r**a**" is valid (i.e., the three A's in this name are a reminder that the AAA argument is a valid one).

Now as much as you would probably love to learn some of the songs and poems used over the years to keep track of Aristotle's system, there is a simpler graphical way to represent syllogistic arguments, called a Venn diagram, that is more suitable for our present purposes.

For example, here is a diagram representing the first premise of our original argument:

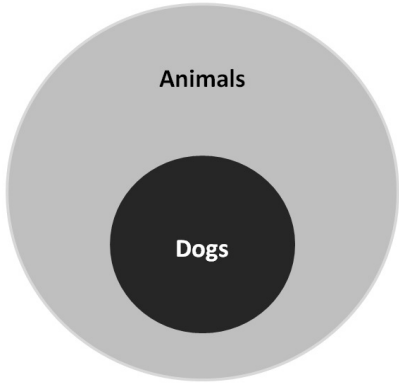

In this Venn diagram, the circle labeled "Dogs" is completely contained within the circle labeled "Animals" so that there cannot be anything that is a dog that is not also an animal. So this diagram illustrates the logical statement "all dogs are animals."

You can do something similar with an E statement such as "No dogs are cats," which would be diagrammed as follows:

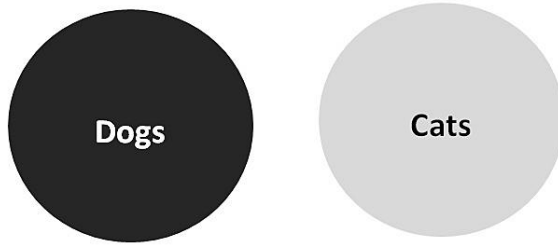

As you can see, the complete lack of overlap between the "Dogs" and "Cats" circles means there is no way for something to simultaneously be a dog and a cat (which would require one thing to fit into the separated "Dogs" and "Cats" groups at the same time, which is impossible).

I and O statements that make use of the word "some" are represented by overlapping circles, as in "Some women are accountants," which would be represented in this manner:

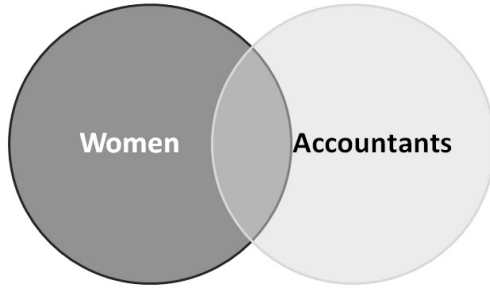

In this case, the overlap between the circles represents the "some" being talked about in the statement. The statement "Some accountants are not women" would be represented by the same overlapping circles shown above, but in that case the entire "Accountants" circle except for the overlap with the "Women" circle would be darkened, indicating that this is the "some" (non-woman accountants) the statement is describing.

While it is nice to use Venn diagrams to illustrate individual components of an argument, they really become useful when you

use them to illustrate the argument as a whole. For example, if you diagram both premises in the animals/dogs/Francine example we've been using, it would look like this:

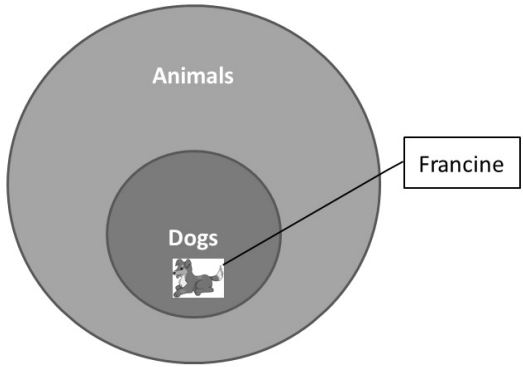

Here, Francine is represented by a picture that is inside the circle labeled "Dogs." As you can see, since the "Dog" circle is completely contained within the "Animals" circle, there is no way that Francine can be a dog (i.e., be placed within the "Dogs" circle) without also being an animal (i.e., automatically being located in the "Animals" circle). Thus, the argument is valid.

We'll be using similar diagrams to illustrate more complex arguments later. But for now, I'd like to introduce a second test, one separate from validity, that an argument has to pass in order for it to be any good: the test of *soundness*.

To understand soundness, let's look at this simple argument:

Premise 1: No mermaids are leprechauns.
Premise 2: Gwen is a mermaid.
Conclusion: Therefore, Gwen is not a leprechaun.

If you diagram this argument, you will see that it is valid, which means there is no way for you to accept the premises but then reject the conclusion.

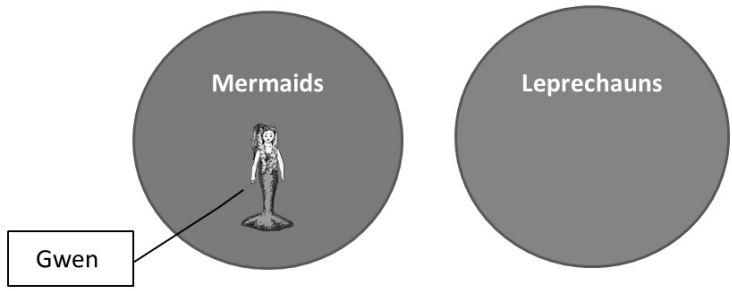

But even if this argument is valid (which it is), why should you accept the premises? After all, here on earth leprechauns and mermaids do not exist, which means that the premises are false (or at least questionable by a reasonable person). And if even one premise in a deductive argument can be proven wrong (or simply doubtful), then the entire argument is considered *unsound*.

Switching back from Gwen the mermaid to Francine the dog, if you look at the premises in that argument, one is clearly correct but the other is up for grabs:

Premise 1: All dogs are animals.
Premise 2: Francine is a dog.
Conclusion: Therefore, Francine is an animal.

If we are using the words "dogs" and "animals" in the normal way they would be used in a sentence like the one in Premise 1, then our first premise ("all dogs are animals") presents no problems.[3] But what about the second premise that I'm asking you to accept ("Francine is a dog")? How well does that one hold up?

[3] In order for these types of logical approaches to work, we need to assume statements reflect standard vs. novel use of language. For example, one could claim "all dogs are animals" is false if "dog" is defined as something of low quality (as "the car he just bought is a real dog"). But while natural language contains ambiguity, when translating human language into logical statements, the goal is to come up with wording that will be unambiguous to any reasonable person.

Well, have you seen Francine or do you have other ways of knowing for sure she's a dog? Or are you just taking my word for it? But I could be lying! Francine could be my car or my kid's stuffed animal. If that were the case, then Premise 2 would be false, which would make the entire argument unsound regardless of its validity.

You should also note that not only can valid arguments be unsound (and thus no good) but that unsound arguments can have true conclusions. For example, if Francine is my cat, then the original argument would still be valid and the conclusion true (since Francine the cat is also an animal). But the entire argument would still be unsound.

Let's jump to a political example, both to break out of the menagerie and to illustrate a few more important concepts related to syllogisms.

Let's say someone at a party tells you "I'd never vote for a Republican since they're all a bunch of warmongers!" Well, in addition to being provocative, they are also making an argument. (Don't worry partisans; I'll be using a similar Democratic example shortly.) Here, the argument (like most arguments carried on by humans vs. those made by lunatics, cyborgs, or logicians) is not written in the syllogistic language of formal logic. But we can do a bit of translation to fit this statement into the proper syllogistic structure.

In this case, the sentence starts with the conclusion that our speaker would never vote for a Republican, after which a premise justifying that conclusion is presented. So if we were to translate the original statement into a syllogistic argument, it would look like this:

Premise 1: All Republicans are warmongers.
Conclusion: Therefore, I would never vote for a Republican.

You'll notice that something is missing that is required for this to be a proper syllogism, namely a second premise. This happens because, in many situations involving normal conversation, a needed premise is implied rather than stated outright. As with other concepts related to logic, there's a Greek word for the concept of a missing premise, in this case an *enthymeme*. The process of translating an argument from real-world language into some kind of formal structure often involves teasing out this missing premise (that is, the enthymeme) so it can be included in the syllogism.

In this example, this currently missing second premise needs to connect the "I" in the conclusion to the "warmongers" term in the first premise. With a little thought, the statement that would accomplish this goal can be added to our argument as follows:

Premise 1: All Republicans are warmongers.
Premise 2: I would never vote for a warmonger.
Conclusion: Therefore, I would never vote for a Republican.

With our syllogism in place, we can now test it for validity and soundness. As you may be able to see by now, the argument is valid in that, if you accept both premises, you have no choice but to accept the conclusion. If you don't believe me, here's a diagram to prove it:

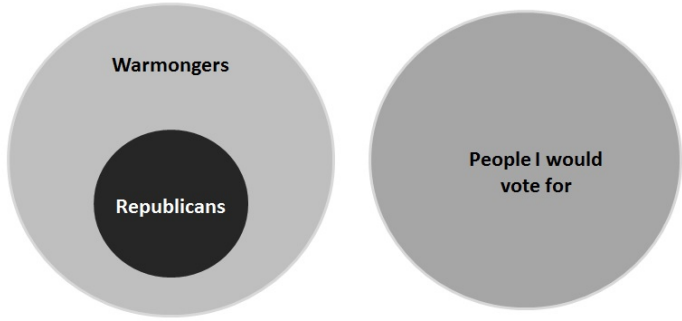

As you can see, there is no way something can be in both the set of warmongers (in which all Republicans are included, at least in this argument) and a second set of people I would vote for. So the argument is valid. But is it sound?

I'm guessing that even the most partisan among you might have problems with the term "warmonger," since, unlike "dogs" and "animals," this word would be a challenge to define in a way everyone (or even many people) would likely agree on. Even so, our new Premise 2 ("I would never vote for a warmonger") is still OK since, regardless of how the arguer defines "warmonger," that would likely describe someone he or she would never vote for.

But the first premise, "all Republicans are warmongers," is saying that however that term is defined, every member of the Republican Party (currently more than fifty million strong) fits into that category. Given the unlikelihood that such a large and diverse group of people would be united by this characteristic (especially since most of them would reject the designation), it's safe to say Premise 1 is a clunker. Since all we need is one dud premise to ruin our argument, then we are left with an argument that, while completely valid, fails due to lack of soundness.

I'd now like to present an argument that offers an opportunity to break out of the straitjacket of Aristotle's syllogisms, one which will also let me keep my promise to ding both parties in my examples. So, let's say it's 1984 and someone at the same party where that last crack about Republicans was made says this as a rejoinder:

"Oh yeah! Well I'd never vote for a Democrat since they're all soft on defense. Which is why I'm not voting for Walter Mondale!"

Translating this argument will be similar to our last example. Ignoring the "Oh yeah!" as superfluous, we can turn our partygoer's response to the following:

Premise 1: All Democrats are soft on defense.
Conclusion: Therefore, I would never vote for Walter Mondale.

Just as in the last example, we can provide a missing premise that would link the "I" in the conclusion to the term "soft on defense" in an explicitly stated premise, which would get us to this:

Premise 1: All Democrats are soft on defense.
Premise 2: I would never vote for someone who is soft on defense.
Conclusion: Therefore, I would never vote for Walter Mondale.

Even with this addition, however, something still seems to be missing, a problem that can be solved if we add just one more premise to our argument as follows:

Premise 1: All Democrats are soft on defense.
Premise 2: I would never vote for someone who is soft on defense.
Premise 3: Walter Mondale is a Democrat.
Conclusion: Therefore, I would never vote for Walter Mondale.

With the addition of a third premise, we have left the realm of formal syllogisms and entered the world of general argumentation in which our arguments can include as many premises as we like (or at least as many premises as are required to state our case).

Just as with a two-premise syllogism, an argument with more than two premises is only considered valid if accepting the premises (all of them) forces you to also accept the conclusion. And the argument is only sound if none of the premises can be proven false or doubtful.

Just like our "Republicans are warmongers" argument that started off this heated exchange, this three-premise argument is

valid (as diagramed below) but unsound since "soft on defense" is too vague and controversial a term to apply to the tens of millions of people who are members of the Democratic Party.

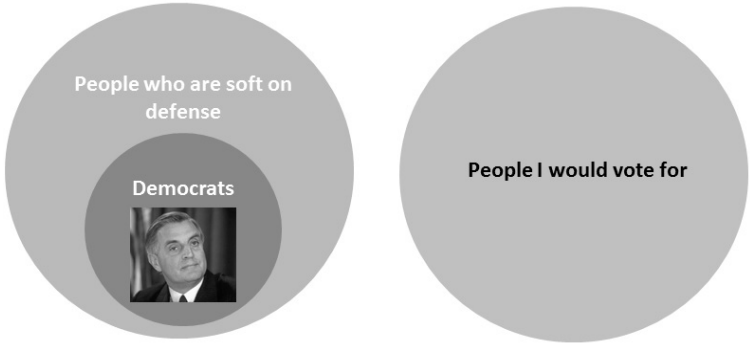

Since many of the arguments that take place in normal human language require more than two premises, this general structure for an argument is more versatile than Aristotle's syllogisms, even if expanding the number of premises means we can no longer use AEIO mnemonics to sort good arguments from bad. But diagraming still works (as you just saw). And the general rules regarding how to determine validity and soundness still apply when an argument includes three or more premises.

In the "Case Studies" section of this book, we'll look at more complex examples of real-world argumentation (like negative ads) that translate to arguments containing three or more premises. But before leaving the topic of arguments that originate in normal human language, I'd like to note two rules that should be applied when translating an argument from normal conversation into any structured format that will be used to analyze that argument for strength and weaknesses.

The first is the *rule of economy*, which asks you to try to keep your number of premises as small as possible while still providing enough of them to support your argument. The reason for this has to do with soundness. After all, it only takes one bad premise (i.e.,

one premise that is false, ambiguous, or can otherwise be rejected by a reasonable person) to take the whole argument down with it. So the more premises you have, the greater the likelihood that one of them will get you into trouble, leaving your entire argument exposed as unsound.

The second is the *rule of charity*, an offshoot of the principle of charity I mentioned in an earlier chapter on bias that requires you to engage with your opponent's strongest arguments rather than just attack his or her weakest ones (or worse, pounce on simple or innocent errors or misstatements not central to the point being made).

When translating arguments from natural to structured language, this principle involves choosing language that accurately reflects what the original arguer was trying to say. For example, in that last argument we analyzed, someone disagreeing with that argument might choose to translate "all Democrats are soft on defense" into something like "all Democrats want to disband the military."

While I already noted that the phrase "soft on defense" is vague (which means that trying to translate it into something more concrete is a legitimate step in the translation process), if you asked the person who originally made the statement whether this translation was acceptable, he or she would likely tell you "No." Which makes it an uncharitable translation, something to be avoided by critical thinkers vs. those who prefer to argue against a point their opponent never made (an unfortunately all-too-common occurrence in our discourse, political or otherwise).

You can probably see how this concept ties into the notion of bias since performing accurate translations requires you to step outside your own biases and get into the head of someone saying things with which you might not agree, all so you can word their arguments in a way an opponent would say accurately reflected their ideas and beliefs. Kevin deLaplante from the Critical Thinker Academy likens the ability to do this work to the skills required to

be a successful actor since critical thinkers should be able to get inside the head of those with whom they disagree to such an extent that they can present an opponent's argument in terms that opponent would say truly reflect what he or she is trying to say.

The translation of arguments is both an art and a science, and you will be exposed to many more examples of this process before you reach the end of this book. But if I can make one last observation before you start practicing this skill on your own: there is often a benefit to putting in the effort needed to translate a real-world argument into a structured format that is valid.

Doing so should not be considered just an act of charity. For most arguments that take place between people usually suffer not from lack of validity but from lack of soundness and forcing an argument into a valid structure tends to require writing premises in such a way that their weaknesses are exposed (as you saw in the valid but unsound arguments we've been using as examples throughout this chapter).

Before closing out this discussion of logic, I'd like to keep an earlier promise by saying a few things about inductive vs. deductive arguments.

As mentioned in the last chapter, inductive logic moves from known facts to reasonable conclusions drawn from those facts. For example, turning that "it's always rained in Massachusetts so it's going to rain this year" argument I mentioned in that chapter into a structured inductive argument would give us:

> **Premise 1**: It's rained every year in Massachusetts in recorded history.
> **Conclusion**: Therefore, it's going to rain again in Massachusetts this year.

In this example, there is no need to look for additional hidden premises since we are just making a perfectly reasonable inference from the one premise that has been presented. In fact, the only

thing that would tighten up this structured inductive argument would be a qualifier in the conclusion (such as "It's really, really likely to rain in Massachusetts this year."), one that clearly identifies that our conclusion reflects high probability vs. absolute certainty.

If you want a more complex (and political) example, consider this conversation taking place around the time of the 2004 presidential election:

"So how's Johnny going to vote? Well, he's a registered Republican. He's voted for someone called George Bush in three presidential elections. And he's working for the Bush campaign. So I'm going to go out on a limb and say that he'll be voting for W."

This argument contains a few rhetorical flourishes (including *litotes* or purposeful understatement, which will be described in a later chapter on rhetorical devices). Stripped of elements that make the conversational version of this argument more entertaining, however, the argument itself looks something like this:

Premise 1: Johnny is a registered Republican.
Premise 2: Johnny voted for George H.W. Bush in two Presidential elections.
Premise 3: Johnny voted for George W. Bush in the last election.
Premise 4: Johnny is working on the campaign to get George W. Bush elected president.
Conclusion: Therefore, Johnny is very likely going to vote for George W. Bush in the next election.

Now there are a couple of things worth noting about this formally structured inductive argument.

First, you'll notice that the general structure we introduced in this chapter for how to organize arguments (with as many premises

as necessary written above a conclusion) works perfectly well for inductive as well as deductive arguments.

That said, a successful inductive argument does not require you to accept the conclusion if you have accepted the premises. In this argument, for example, it's certainly possible that Johnny, despite his voting record and current volunteer activities, pulled the lever for John Kerry when he reached the voting booth in 2004 (either by accident or after discovering that George W. Bush was a disguised alien, for example). Possible but unlikely, which is why the conclusion you are being asked to accept just requires you to embrace high probability vs. metaphysical certainty.

Also, unlike deductive arguments, inductive arguments aren't ruined by one false or otherwise crappy premise. For example, what would be the impact on this argument if it turns out that the speaker was wrong and that Johnny is registered as an Independent? While realizing Premise 1 is false in such as way would certainly *weaken* the argument, the other premises still provide more than enough evidence to support the conclusion.

Given that more premises are better than fewer with regard to inductive arguments, the rule of economy might not seem to apply since any evidence that increases the likelihood that a conclusion is true can strengthen an argument. Interestingly enough, though, piling on too many premises can actually make an argument *seem* weaker.

For instance, we could add to our 2004 election argument the fact that Johnny has a Bush bumper sticker on his car and that his parents are big donors to the Republican Party. While both additional premises add to the likelihood that the conclusion is true, their inclusion can seem redundant (why wouldn't a Bush campaign volunteer have a Bush bumper sticker on his car, after all?) or of questionable relevance (since kids rebel against their parents' political beliefs as often as they embrace them).

Like an overzealous prosecutor charging a defendant with every imaginable crime large and small hoping something will stick,

someone packing his or her argument with more and more evidence (particularly of varying quality) can come off as desperate, causing people to judge an argument based on its weakest vs. strongest premises. Which is why you should still think economically when making any kind of argument, inductive or deductive, building it on premises that include just your strongest and most pertinent evidence.

As a final, final note, in contrast to the complex way the rule of economy applies to inductive reasoning, the rule of charity should remain intact since even inductive arguments need to honestly reflect what people are saying.

With that said, it's time to take a closer look at a subject we've been dancing around for several chapters now: argumentation, the primary way we humans communicate, convince, and (ideally) critically reason our way through life.