

Chapter 3

Logical Arguments

Monty Python’s Flying Circus, a late-1960s/early-1970s British sketch-comedy television show, was created by a group of five Oxford- and Cambridge-educated comedy writers and one American animator. Several of the sketches from that show have acquired comedic immortality, including the “Argument Clinic” from the episode “The Money Programme” that aired in November of 1972. Figure 3.1 is a screen-capture. In the bit, a man (“Man,” played by Sir Michael Palin¹) pays for what turns out to be an unsatisfying argument with a second man (“Mr. Vibrating,” John Cleese²). If you have never seen it, you should stop reading this and go watch it.³

Welcome back! In the middle of the non-argument argument, Palin gives a very good definition, which we adopt here.

Definition 16: Argument

“An *argument* is a connected series of statements to establish a definite proposition.”

argument

Our demonstrations of logical equivalence from Chapter 1 are covered by this definition, as we used a sequence of logical expressions to establish that all expressions in the sequence are logically equivalent. Arguments can be much

¹Palin was appointed a Knight Commander of the Most Distinguished Order of Saint Michael and Saint George in 2019, for services to travel, culture and geography.

²Because I know you’re wondering: Cleese has routinely declined royal honors.

³ Many of Monty Python’s sketches, including the Argument Clinic, can be found on YouTube: <https://www.youtube.com/user/MontyPython>



Figure 3.1: Michael Palin, left, argues arguments with John Cleese in Monty Python’s “Argument Clinic” sketch. Credit: BBC.

more involved. This chapter introduces common argument structures and shows how logical equivalences can be combined with rules of logical inference to create complex arguments, our final stop on the road to proofs.

Please be aware that what are often called ‘arguments’ on television and radio talk shows rarely possess a logical foundation. This chapter’s arguments will be based on logic, will not cover politics, and will require no raised voices or finger-pointing.

3.1 Reasoning

Have you ever written down an answer to an exam question for no better reason than it felt like the correct answer? If so, you relied on your *intuition*, for better or worse. You probably feel more confident in answers that you can convincingly justify. The process of drawing justifiable conclusions from given information using principles of logic is *reasoning*.

There are many forms of reasoning, some more formal than others. Here, we have need of two forms of reasoning: *deductive* and *inductive*.

reasoning

3.1.1 Deductive Reasoning

One night, before you go to bed, you set your phone on the nightstand, fully expecting that it will still be there when you wake up in the morning. You have that expectation because you have two facts: You placed the phone on the nightstand, and Newton’s First Law of Motion (a motionless object will remain motionless until a force is applied).⁴ Assuming that nature, pets, and practical-joking roommates fail to apply adequate forces, those two facts will convince you that your phone will be on the nightstand in the morning. The logical connection between those facts and your conviction is *deduction*.

deduction

Definition 17: Deductive Argument

A *deductive argument* uses general principles of logic, applied to a given set of facts, to draw a conclusion from those facts.

deductive argument

Example 55:

Problem: You have a dog as a pet. You have dropped a piece of bread on the floor. You conclude that you won’t have to pick up the piece of bread. Is that a deductive argument?

Solution: Yes, that is a deductive argument. Anyone who has ever had a dog knows that dogs will eat any food – or anything that remotely mimics food – that reaches the ground. Combine that knowledge with the given facts, and you know that the bread is as good as gone.

Example 55’s solution is incomplete – we don’t (yet!) know the general principle of logic that is being applied to those facts to reach that conclusion. Logical justifications are an essential component of a complete deductive argument.

All of the proof forms presented in this book will be forms of deductive arguments . . . even the ones named after induction.

⁴Second definition: An unguarded Fig Newton will move into the nearest mouth.

3.1.2 Inductive Reasoning

induction

You’ve probably noticed that the sun rose in the east two days ago. It did the same thing yesterday morning, and again this morning. Based on that, you probably expect it to rise in the east again tomorrow morning, and every day thereafter.⁵ This is an example of reasoning by *induction* : You started with some observations, and formed a general conclusion based on those observations.

inductive argument

Definition 18: Inductive Argument

An *inductive argument* reaches a general conclusion from a set of specific observations.

Example 56:

Episodes of the Fox television show “Bones” usually start with unsuspecting people finding human remains, and conclude with the perpetrator being captured by an FBI agent and a forensic anthropologist. The following dialog appears in the first-season episode “Woman in the Tunnel” as the team examines a virtual representation of tunnels under a city:

⁵Until the Earth’s magnetic poles flip and the sun starts rising in the west. That won’t happen for thousands of years . . . probably.

Booth: Well, we found that Civil War victim near a cave-in. Maybe the treasure's on the other side?

Goodman: Inductive, reductive or deductive?

Brennan: Deductive.

Goodman: As you wish. Ms. Montenegro, please remove all tunnels containing power, cable or utility lines.
[...]

Booth: Oh, what about diamond dust? You said that there was diamond dust in the old tunnels. There was also diamond dust on the Civil War guy, so ... what? I'm not allowed to help now?

Goodman: That's inductive logic.

Brennan: We agreed on deductive.

Booth: I'm sorry, I'm just, you know, trying to think outside your box.

The Goodman character deduces from the facts that the body was from the U.S. Civil War period and that people in the 1860s did not have amenities such as cable television to decide that a body from that time could not have been resting in a tunnel with modern wiring. Later, Booth tries to generalize from specific examples of the appearance of diamond dust to draw a conclusion. That's inductive reasoning.

The dialog used in Example 56 includes a reference to *reductive* reasoning, which is also a useful if somewhat confusing way to form an argument. We won't discuss it in this chapter, but we will in Chapter 4.

Defining inductive arguments isn't as straight-forward as defining deductive arguments. In philosophy, inductive reasoning allows for uncertainty – an inductive argument strongly suggests that the conclusion is correct, but doesn't claim that it is correct. In mathematics and the sciences, we crave certainty, but the structure of inductive reasoning is appealing because it is directly applicable to a variety of provable situations. To address this dilemma, *mathematical induction* wraps deductive reasoning in a shell of inductive reasoning, giving us the form of induction but the certainty of deduction. We will cover mathematical induction in a later chapter.

3.1.3 Abductive Reasoning

The American city of Tucson, Arizona appears to be roughly surrounded by mountain ranges, including the Rincons to the east, the Santa Catalinas to



Figure 3.2: Mountain ranges near Tucson, Arizona. Credit: Arizona Daily Star.

the northeast, the Tortolitas to the northwest, and the Tucson Mountains to the west (see Figure 3.2). It's tempting to assume that this current geologic arrangement was formed from the collapse of an ancient and massive volcano's caldera, with the current mountain ranges being the uncollapsed edges of the volcano.⁶

abduction

Such an assumption can arise from abductive reasoning. In *abduction*, we start with the end result and propose possible situations that could have produced that result. This is roughly the converse of the deductive reasoning process combined with the uncertainty of inductive reasoning. As such, it isn't useful logically, although it can be helpful for brainstorming ideas that might turn out to be provable.

Example 57:

Problem: In Gotham City, Detective Montoya is investigating a homicide.

⁶Tempting, but likely incorrect. Based on the types of rocks, the current thinking is that the volcano was where the Catalinas are, and detachment faults caused the Tucson Mountains to 'slide' to the southwest. The city is built on the mile-deep sediment that rests in-between.

While searching the victim, she finds a joker playing card in his hand. Without hesitating, she pulls out her phone and calls Commissioner Gordon. Why did she feel it necessary to ignore her chain of command and call him directly?

Solution: Through abductive reasoning, of course! In Gotham, a dead man holding a joker could mean that the Joker has escaped from Arkam Asylum and has killed again. Or, perhaps the Penguin is the real culprit and is trying to shift the blame. Either way, it's time to call Batman. The brainstorming that occurs when using abductive reasoning is helpful when investigating crimes.⁷

Given its speculative nature, we will consider abduction no further.

3.2 Categorizing Deductive Arguments

For the rest of this chapter, we will consider only deductive arguments. Some deductive arguments are more speculative than others. In this section we will discuss that distinction.

3.2.1 Writing Deductive Arguments

In section 1.3.6, we learned that in an implication, the left operand can be called the hypothesis and the right the conclusion. Deductive arguments can be viewed as implications with multiple hypotheses (some given to us and some deduced by us) and a single ultimate conclusion.

Example 58:

We can rewrite the deductive argument from Example 55 as an if-then sentence: *If you have a dog and you dropped a piece of bread on the floor, then you did not need to pick up that piece of bread.* This makes the individual hypotheses of the compound hypothesis, and the intended conclusion, clear. Written in logic, we have $(d \wedge b) \rightarrow \neg n$, where d is 'you have a dog', b is 'you dropped a piece of bread on the floor', and n is 'you need to pick up that piece of bread'.

⁷Even fictional ones. Yes, Batman is a fictional character. Deal with it.

Although the hypotheses and the conclusion are clear, the justification that allows us to reach that conclusion from those hypotheses is not. As mentioned earlier in this chapter, a complete deductive argument also includes justifications for its steps. To make it easy for us to include these explanations, we will adopt a vertical format for our argument components. For example, here's the vertical form of the logical expression $(d \wedge b) \rightarrow \neg n$ from Example 58:

$$\frac{\begin{array}{c} d \\ b \end{array}}{\therefore \neg n}$$

The vinculum⁸ separates the hypotheses from the conclusion, and the hypotheses are assumed to be ANDed together. The symbol \therefore (`\LATEX : \therefore`) means ‘therefore’ and is used to highlight the conclusion. For such a short expression, we could adopt the horizontal equivalent: $d, b / \therefore \neg n$. The advantage of the vertical form is that we have room to the side to provide each line with a justification. While we're at it, we can also number the lines, to make referencing them easy:

$$\begin{array}{lcl} (1) & d & [\text{Given}] \\ (2) & b & [\text{Given}] \\ \hline (3) & \therefore \neg n & [\text{Ummmm ... magic?}] \end{array}$$

Sadly, we still don't know how to get to the conclusion. But we will!

3.2.2 Valid and Sound Arguments

For an argument to be convincing, we must be able to see that the truth of the conclusion is based on the truth of the available hypotheses. Such an argument is said to be *valid*.

valid argument

Definition 19: Valid Argument

Any deductive argument of the form $(p_1 \wedge p_2 \wedge \dots \wedge p_n) \rightarrow q$ is *valid* if the truth of conclusion q can be derived from the truth of hypotheses p_1 through p_n .

⁸The horizontal line, but you already knew that because you've read the math review appendix ... right?

What the definition of a valid argument doesn't say is whether or not the hypotheses are actually true! This might appear to be a huge problem, but it isn't. If you think back to abductive reasoning for a moment, you'll recall that it was speculative – we thought of one or more situations that could have caused the conclusion to result. We can assume the truth of the hypotheses that describe one of those situations, correctly apply logic to those hypotheses, and determine if the conclusion follows. If it does, the validity of the argument has been logically established, even though the hypotheses are not yet known to be true.

Example 59:

Consider the following argument:

- | | | |
|-----|--|--|
| (1) | 63 is evenly divisible by 2 | |
| (2) | 63 is evenly divisible by 3 | |
| | | |
| (3) | \therefore 63 is evenly divisible by 6 | |

What do you think of that argument? If you are dismissing it because 63 is not evenly divisible by two, you're missing the point. The argument itself is valid: When you have a number that is divisible by both two and three, it will also be divisible by six. Remember, in a valid argument, the hypotheses need only be assumed true.

Valid arguments, then, have room for improvement: We could insist that the hypotheses actually be true! For the argument to be anything more than an exercise in logical construction, they need to be. If they are, we have more than a valid argument – we have a *sound* one.

Definition 20: Sound Argument

Any valid deductive argument is *sound* when its hypotheses are true.

sound argument

Because 63 is not actually evenly divisible by two, the argument of Example 59 is not sound. Changing 63 to any multiple of six will make it both valid and sound. Please note that a sound argument is necessarily valid, just as a multiple of six is necessarily a multiple of two.

3.3 Rules of Inference for Propositions

Let's reflect on what a valid argument tells us: When the hypotheses are true, the conclusion is true. Notice that a valid argument tells us nothing about the truth of the conclusion when one or more of the hypotheses are false.

Because the hypotheses of a valid argument supply enough information to reach the conclusion, we can demonstrate that validity with a truth table.

Example 60:

Consider the following trivial but valid argument:

$$\begin{array}{l} (1) \quad p \\ (2) \quad q \\ \hline (3) \quad \therefore p \wedge q \end{array}$$

We can rewrite this argument in the form of the logical expression $(p \wedge q) \rightarrow (p \wedge q)$. We don't need to create a truth table to see that this expression is true when both p and q are true, but we'll do it anyway to make a point:

p	q	$p \wedge q$	$(p \wedge q) \rightarrow (p \wedge q)$
T	T	T	T
T	F	F	T
F	T	F	T
F	F	F	T

First, look at the third column (labeled $p \wedge q$), which represents our conclusion. When both of the hypotheses are true (the first row), the conclusion is true. This shows that the argument $p, q / \therefore p \wedge q$ is valid.

Now consider the fourth column, which represents the entire argument. It's a tautology! In this example, that's also not much of a surprise; $p \rightarrow p \equiv \mathbf{T}$ is one of the logical equivalences ("Self-implication") presented in Chapter 1. Still, keep this observation in mind; we will consider it again later in this section and also in section 3.5.1.

Using truth tables to verify the validity of arguments is possible, as we just showed, but is tedious for all but the smallest of arguments. Fortunately, we have an alternative.

Remember, back in Chapter 1, when we used a given⁹ set of logical equivalences to show that more complex expressions were also equivalences? We can do the same sort of thing with arguments: We can construct large arguments using smaller arguments as building blocks. We will also make use of those logical equivalences to restate expressions into forms that suit our needs.

3.3.1 Eight Rules of Inference

The tiny valid argument used in Example 60 is known as a *rule of inference*. Many such tiny valid arguments exist, but only a few of them are used often enough in this book to be worth learning. Some, such as the one used in Example 60, are clearly valid. Others can be shown to be, again using a truth table.

rule of inference

Table 14: Some Rules of Inference Worth Knowing
--

⁹But verifiable – you don’t have to trust us!

	Name (Alternate Name)	Argument
(1)	Conjunction (Conjunction Introduction)	$\frac{p}{q}$ $\therefore p \wedge q$
(2)	Simplification (Conjunction Elimination)	$\frac{p \wedge q}{p}$ $\therefore p$
(3)	Addition (Disjunction Introduction)	$\frac{p}{p \vee q}$ $\therefore p \vee q$
(4)	Disjunctive Syllogism (Disjunction Elimination)	$\frac{p \vee q}{\bar{p}}$ $\therefore q$
(5)	Modus Ponens (Affirming the Antecedent)	$\frac{p}{p \rightarrow q}$ $\therefore q$
(6)	Modus Tollens (Denying the Consequent)	$\frac{\bar{q}}{p \rightarrow q}$ $\therefore \bar{p}$
(7)	Hypothetical Syllogism (Transitivity of Implication)	$\frac{p \rightarrow q}{q \rightarrow r}$ $\therefore p \rightarrow r$
(8)	Resolution	$\frac{p \vee q}{\bar{p} \vee r}$ $\therefore q \vee r$

syllogism

The word *syllogism* is a synonym for deductive reasoning, making its use in the name of an argument a bit redundant. But, those are the names most people use, and so shall we.

If you take the time to verify the validity of these rules, you will probably notice that they are all tautologies, as we'd hinted at in Example 60. This isn't coincidence. Recall that arguments (including rules of inference) have the form $(p_1 \wedge p_2 \wedge \dots \wedge p_n) \rightarrow q$. We know that, for an argument to be valid, when the hypotheses are all true, the conclusion *must* be true. In this case, the complete argument will also be true $((\mathbf{T} \wedge \mathbf{T} \wedge \dots \wedge \mathbf{T}) \rightarrow \mathbf{T} \equiv \mathbf{T} \rightarrow \mathbf{T} \equiv \mathbf{T})$. But even if one or more of the hypotheses are false, the argument will still be true $((\mathbf{T} \wedge \mathbf{F} \wedge \dots \wedge \mathbf{T}) \rightarrow \textit{whatever} \equiv \mathbf{F} \rightarrow \textit{whatever} \equiv \mathbf{T})$. Thus, valid

arguments (including rules of inference) must always evaluate to true; in other words, they are always tautologies.

Occasionally, people lump rules of inference together with logical equivalences and attempt to use the former as if they were the latter. Rules of inferences do **not** claim that the hypotheses are logically equivalent to the conclusions. That may sometimes be true (as it is with Conjunction, for example), but most often it is not.

3.3.2 Using Rules of Inference in Arguments

We know about logical equivalences, valid arguments, and rules of inference. It's time to put them to work producing conclusions from given hypotheses. We will start with a simple argument.

Example 61:

Problem: Given $\neg b$ and $a \rightarrow b$, what can we conclude about a ?

Solution: By modus tollens, we can conclude that a must be false. Here's the complete argument, in our preferred vertical form:

$$\begin{array}{lll} (1) & \neg b & [\text{Given}] \\ (2) & a \rightarrow b & [\text{Given}] \\ \hline (3) & \therefore \neg a & [1, 2, \text{Modus Tollens}] \end{array}$$

By listing $\neg b$ as a hypothesis, we are assuming it to be true. For the negation of b to be true, b itself must be false. This is how we can say something is false when every expression is assumed true. In the same way, the conclusion of $\neg a$ says that we have determined a to be false.

Notice that we justified the conclusion by including hypothesis line numbers with the rule of inference we used. We recommend that you adopt this practice, too, because doing so makes it easier for the reader to follow the logic of the argument. In a small example such as this, it doesn't really help, but in longer arguments it will.

The next example is almost as simple, but we need to perform some con-

versions from English to Logic.¹⁰

Example 62:

Problem: You are ready to go to bed, but see some light coming from the other end of the hallway. That means that you forgot to turn off either the dining room light or the kitchen light. You stick your head out the bedroom door and see that the kitchen is dark – the kitchen light is off. How are you able to conclude that the dining room light is on?

Solution: It’s tempting to answer, “Because the kitchen light’s not on, duh!” While it’s true that the kitchen light isn’t on, that fact alone does not explain the source of the light. We need to construct a valid and completely justified argument from the given hypotheses, showing that the conclusion is true if the hypotheses are.

First step: Identify the hypotheses and conclusion. In other words, what do we know and what do we hope to show? We know that the dining room light is on or the kitchen light is on. We also know that the kitchen light is off. We hope to show that the dining room light is on. These statements will be easier to work with if we give their component propositions labels, remembering that the fact that the kitchen light is off is not a third proposition, but rather the negation of the second:

d : The dining room light is on
 k : The kitchen light is on

We can re-write our hypotheses and conclusion as logical expressions using these labels:

Hypotheses:

“the dining room light is on or the kitchen light is on” $\Rightarrow d \vee k$
 “the kitchen light is off” $\Rightarrow \neg k$

Conclusion:

“the dining room light is on” $\Rightarrow d$

How can we show that the conclusion follows from those hypotheses?

¹⁰Bet you wished you’d never have to do that again! Like it or not, it’s a necessary step in many arguments, not to mention proofs.

Table 14 includes a rule of inference that fits our situation: Disjunctive Syllogism. In fact, it fits so well that it completes the argument all by itself. All we have to do is write out the argument with our adornments:

(1)	$d \vee k$	[Given]
(2)	$\neg k$	[Given]
(3)	$\therefore d$	[1, 2, Disjunctive Syllogism]

If you are particularly attentive, you may be concerned that the order of the labels in our first proposition is the opposite of the order used in the form of disjunctive syllogism in Table 14. That is, in the table, we have $p \vee q$ and $\neg p$, meaning that the negation of the left operand is true. In the example, we have $d \vee k$ and $\neg k$, where the negation of the right operand is true. This really isn't a problem. Disjunctive syllogism says that if we have an inclusive disjunction and one of the operands of this disjunctions is false, the other must be true. It doesn't matter which one is false, because inclusive-OR is commutative. If this implicit step bothers you, you can add a line to the argument that commutes the disjunction (and uses 'Commutativity of Disjunction' to justify it) before applying disjunctive syllogism.

Similarly, we could have listed $\neg k$ as line (1) and $d \vee k$ as line (2) without causing any trouble. Disjunctive syllogism would still be applicable because the hypotheses are being ANDed together (see Definition 19), and conjunction is also commutative.

Example 63:

Problem: What is wrong with this valid argument?

If seven is larger than eight, then seven minus eight is a positive number. Seven is larger than eight. Therefore, seven minus eight is a positive number.

Solution: There's nothing wrong ... with the argument! It's a perfectly lovely valid argument, as the following demonstrates:

s : Seven is larger than eight
 m : Seven minus eight is a positive number

(1)	$s \rightarrow m$	[Given]
(2)	s	[Given]
(3) $\therefore m$		
		[1, 2, Modus Ponens]

If we assume the hypotheses to be true, the conclusion logically follows, thanks to modus ponens. If you were thinking that s 's offensive lack of truth invalidated the argument, you were thinking about a sound argument, where the hypotheses need to be true. In valid arguments, we can indulge flights of fancy.

Ready for a more complex argument?

Example 64:

Problem: Your car is parked facing the sun. You do not have a windshield sun shade. When your car is facing the sun without a sun shade, the steering wheel is in the sun. The steering wheel is hot, or it is not in the sun. Is the steering wheel hot?

Solution: It certainly seems reasonable that the steering wheel is hot, but we need to show how to reach that conclusion from the given information and our knowledge of logic.

As always, the first step is to identify and label the propositions, remembering to stay positive so as not to hide any negations:

- f : The car is facing the sun
- s : The car has a sun shade
- w : The steering wheel is in the sun
- h : The steering wheel is hot

Next, express the hypotheses and conclusion in logic notation:

The car is facing the sun	\Rightarrow	f
The car has no sun shade	\Rightarrow	$\neg s$
When the car faces the sun without a sun shade, the steering wheel is in the sun	\Rightarrow	$(f \wedge \neg s) \rightarrow w$
The steering wheel is hot or is not in the sun.	\Rightarrow	$h \vee \neg w$
The steering wheel is hot	\Rightarrow	h

Now the fun starts: We need to see what new truths we can uncover from the hypotheses; hopefully, h will be one of them. One way to approach this task is to start with what we're given and see where it leads us. Hopefully, it will lead us close enough to the conclusion that we'll be able to see how to complete the argument.

Looking at the third hypothesis, we can see how modus ponens can be used to establish the truth of w , if we knew that $f \wedge \neg s$ were true. Ah, but we can show that, from the first two hypotheses!

(1)	f	[Given]
(2)	$\neg s$	[Given]
(3)	$f \wedge \neg s$	[1, 2, Conjunction]
(4)	$(f \wedge \neg s) \rightarrow w$	[Given]
(5)	w	[3, 4, Modus Ponens]

Great ... but does knowing the truth of w help? The fourth hypothesis includes a $\neg w$; can that help? It almost fits the form of disjunctive syllogism, but not quite. Maybe we can find a logical equivalence that can help ... Ah! The 'something ORed with the negation of something else' form appears in the Law of Implication:

(6)	$h \vee \neg w$	[Given]
(7)	$\neg w \vee h$	[6, Commutativity of \vee]
(8)	$w \rightarrow h$	[7, Law of Implication]

(Again, many people would not include line 7 in the argument, but you can't go wrong by including it.)

At this point, you can probably see how to finish it off: Applying modus ponens to lines 5 and 8 shows that h is true. Here's the complete argument:

(1)	f	[Given]
(2)	$\neg s$	[Given]
(3)	$f \wedge \neg s$	[1, 2, Conjunction]
(4)	$(f \wedge \neg s) \rightarrow w$	[Given]
(5)	w	[3, 4, Modus Ponens]
(6)	$h \vee \neg w$	[Given]
(7)	$\neg w \vee h$	[6, Commutative Laws]
(8)	$w \rightarrow h$	[7, Law of Implication]
(9)	$\therefore h$	[5, 8, Modus Ponens]

Yes, the steering wheel hot, and we have explained why.

There's lots to say about Example 64:

1. A complete argument is frequently constructed of many small arguments. Here, we built our argument from two rules of inference and two logical equivalences.
2. Each use of a rule of inference or a logical equivalence created a new (derived) hypothesis that helped move us to our goal. Remember, both logical equivalences and rules of inference produce an expression that we can assume to be true.
3. Numbering the lines of the argument helps!
4. We used modus ponens twice in that one argument. Re-using rules of inference and logical equivalences in the same argument is fine.
5. Some people like to list all of the givens at the start of the argument (i.e., the g givens appear on lines 1 through g), while others prefer to write them into the argument just before they are needed (as we demonstrated here). Either approach is fine.

6. It's not uncommon to reason your way into a dead-end – an expression that doesn't show any sign of helping you reach your conclusion. Don't panic! Back up a step or two, and look through your rules of inference and logical equivalences again. You might find another option that will turn out to be more helpful.
7. Not every desired conclusion can be reached from every initial set of hypotheses. Example arguments in textbooks are usually (though not always) constructed to work.

The argument in Example 64 used a “top-to-bottom” technique. We can also build arguments “bottom-to-top,” as the next example demonstrates on the same set of hypotheses and the same conclusion.

Example 65:

Problem: Can you construct a different argument based on the hypotheses of Example 64 that reaches the same conclusion?

Solution: Definitely! The more complex the argument, the more ways there will be to reach the conclusion (assuming that reaching the conclusion is possible, of course).

We have the same hypotheses and desired conclusion, repeated here for convenience:

Hypotheses:	f
	$\neg s$
	$(f \wedge \neg s) \rightarrow w$
	$h \vee \neg w$
Conclusion: h	

This time, let's start at the conclusion. Can we see a way to reach h from the hypotheses? h appears only in the fourth hypothesis. We'd noted in Example 64 that disjunctive syllogism didn't quite fit that hypothesis, but it's not hard to make it fit. The structure of disjunctive syllogism needs one of the operands in the disjunction to be ‘combined’ with its negation. As we need to keep h , we need to combine $\neg w$ with its negation. The negation of $\neg w$ is $\neg\neg w$, which (by double negation)

is equivalent to w . Thus, if we can show that w is true, we can apply double negation and disjunctive syllogism to show the truth of h . This is how the argument will end:

$(n - 4)$	w	[we hope!]
$(n - 3)$	$\neg\neg w$	[$n - 4$, Double Negation]
$(n - 2)$	$h \vee \neg w$	[Given]
$(n - 1)$	$\neg w \vee h$	[$n - 2$, Commutative Laws]
(n)	$\therefore h$	[$n - 3, n - 1$, Disjunctive Syllogism]

Now, how can we establish the truth of w ? Well, why not the same way we did it in Example 64? We could, but we've been there and done that – let's be different. We can apply the law of implication to $(f \wedge \neg s) \rightarrow w$, producing $\neg(f \wedge \neg s) \vee w$. We already know $f \wedge \neg s$ is true. This is set up perfectly for disjunctive syllogism, giving us a final argument that is almost completely different than that of Example 64:

(1)	$(f \wedge \neg s) \rightarrow w$	[Given]
(2)	$\neg(f \wedge \neg s) \vee w$	[1, Law of Implication]
(3)	f	[Given]
(4)	$\neg s$	[Given]
(5)	$f \wedge \neg s$	[3, 4, Conjunction]
(6)	$\neg\neg(f \wedge \neg s)$	[5, Double Negation]
(7)	w	[2, 6, Disjunctive Syllogism]
(8)	$\neg\neg w$	[7, Double Negation]
(9)	$h \vee \neg w$	[Given]
(10)	$\neg w \vee h$	[9, Commutative Laws]
(11)	$\therefore h$	[8, 10, Disjunctive Syllogism]

The final argument is two steps longer than the one developed in Example 64, but is just as correct, and correctness is more important than brevity.

We've already mentioned that people often leave out commutativity as a step in arguments, assuming that the concept is so basic that it doesn't need to be stated. People often assume the same with double negation. Even if

you're not required to include them, remember that your arguments will be more complete – and easier for people to follow – with those steps in them.

To finish this section, let's fulfill a promise by (finally!) completing the argument introduced in Examples 55 and 58.

Example 66:

Problem: Which additional hypotheses must be assumed true to conclude that you don't need to pick up a piece of dropped bread, assuming that if you have a dog and you dropped the piece of bread on the floor, you don't need to pick it up?

Solution: If you followed the recent argument examples, this should be easy. Here are the labels introduced in Example 58:

- d : You have a dog
- b : You dropped a piece of bread on the floor
- n : You need to pick up that piece of bread

To reach the conclusion ($\neg n$) from the one known hypothesis ($(d \wedge b) \rightarrow \neg n$), we can use modus ponens ... if we could assume that $d \wedge b$ is true. Thus, the answer to the question is: The hypotheses d and b must both be assumed to be true. Here's the complete argument:

(1)	d	[Given]
(2)	b	[Given]
(3)	$d \wedge b$	[1, 2, Conjunction]
(4)	$(d \wedge b) \rightarrow \neg n$	[Given]
(5)	$\therefore \neg n$	[3, 4, Modus Ponens]

And now you know why that puddle of slobber is in the middle of your kitchen floor and a very attentive dog is sitting next to you with a hopeful look on its face.

3.4 Rules of Inference for Predicates

The eight rules of inference from the last section can be applied to predicates just as easily as they can be applied to propositions. The following example's

argument is that of Example 62, but using predicates instead of propositions.

Example 67:

Problem: Restate the argument from Example 62 using predicates instead of propositions.

Solution: Converting propositions to predicates isn't too difficult. The general concept expressed by the proposition becomes the predicate, and the 'subject' is generalized to be the variable. In Example 62, there's just one concept, that of a light being on. Thus, our single predicate is:

$$L(x) : x \text{ is on, } x \in \text{Lights}$$

The argument involves two specific lights: Those of the dining room and the kitchen. Those lights will be constants passed to the predicate to be used to express the hypotheses and conclusion. Here's the re-written argument:

(1)	$L(\text{dining room}) \vee L(\text{kitchen})$	[Given]
(2)	$\neg L(\text{kitchen})$	[Given]
(3) $\therefore L(\text{dining room})$		
		[1, 2, Disjunctive Syllogism]

This application of disjunctive syllogism is correct because “ $L(\text{dining room})$ ” and “the dining room light is on” are just two ways of supplying the same information; that is, two ways of saying the same thing. We think that using predicates when they aren't needed (as in this example) is a bad idea. The predicates add a layer of abstraction and additional syntactic clutter without supplying any benefit. Don't use predicates unless you have to.

3.4.1 Four Rules of Inference for Quantified Expressions

When might you need predicates in an argument? We introduced predicates in Chapter 2 because quantifiers apply to variables and predicates accept variables. Should we need to employ a rule of inference that depends on a

quantifier, we won't have any choice but to express the argument in terms of predicates.

Such rules of inference are not hypothetical. There are four such rules for quantified expressions, two per quantifier, that you may find useful in arguments. Table 15 lists them.

Table 15: Rules of Inference for Quantified Expressions		
	Name	Argument
(1)	Universal Instantiation	$\frac{\forall x P(x), x \in D}{\therefore P(d), d \in D}$
(2)	Universal Generalization	$\frac{P(d) \text{ for every } d \in D}{\therefore \forall x P(x), x \in D}$
(3)	Existential Instantiation	$\frac{\exists x P(x), x \in D}{\therefore P(d) \text{ for some } d \in D}$
(4)	Existential Generalization	$\frac{P(d) \text{ for some } d \in D}{\therefore \exists x P(x), x \in D}$

3.4.2 Using Quantified Expressions in Arguments

Example 68:

Problem: Bicycle Repair Man is Mr. F. G. Superman's secret identity.¹¹ Mr. Superman is a Briton. Does there exist at least one Bicycle Repair Man who is also a Briton?

Solution: It seems very likely that the answer is 'yes,' but how do we know? Is this the right place for an argument?¹²

Our desired conclusion is an existentially-quantified expression, so we need to use predicates to construct the argument. Let's use these ...

$R(x)$: x is Bicycle Repair Man, $x \in \text{People}$

$B(x)$: x is a Briton, $x \in \text{People}$

... to express our hypotheses ($R(\text{F. G. Superman})$ and $B(\text{F. G. Superman})$) and desired conclusion ($\exists x(R(x) \wedge B(x))$). With all of that in place, the argument itself is straight-forward, thanks to existential generalization. Note that we've shortened Mr. Superman's name to get the table rows to be of reasonable lengths:

(1)	$R(\text{F. G. S.})$	[Given]
(2)	$B(\text{F. G. S.})$	[Given]
(3)	$R(\text{F. G. S.}) \wedge B(\text{F. G. S.})$	[1, 2, Conjunction]
(4)	$\therefore \exists x(R(x) \wedge B(x)), x \in \text{People}$	[3, Existential Gen.]

Yes, there is a Briton who is Bicycle Repair Man, to the relief of bicyclists all across England.

You might be wondering why we didn't have a binary predicate (e.g., " $S(x, y)$: The secret identity of x is y "), instead of $R(x)$, so that we could represent the secret identities of many people. We certainly could have, but as we only needed to worry about one superhero, we didn't a more flexible predicate.

Example 69:

Problem: If Alice knows how to swim, so does Bob, but Bob doesn't know how to swim. Create an argument demonstrating that the statement "Someone knows how to swim" is false.

Solution: The first word ('someone') of the desired conclusion makes it clear that we need a predicate covering the ability of a person to swim. Less clear is how far-reaching 'someone' should be – all people or just Alice and Bob? As it's quite a jump from Alice and Bob to every human,

¹¹See the Bicycle Repair Man sketch from Episode 3 ("How to Recognise Different Types of Trees From Quite a Long Way Away") of Monty Python's Flying Circus.

¹²"I've told you once." ☺

we'll assume the set of size two. As a convenience, we use D to label the domain set $\{\text{Alice}, \text{Bob}\}$.

We need just one predicate with which to express our two hypotheses and single conclusion:

Predicate: $S(x) : x$ knows how to swim, $x \in D$

Hypotheses: $S(\text{Alice}) \rightarrow S(\text{Bob})$
 $\neg S(\text{Bob})$

Conclusion: $\neg \exists x(S(x)), x \in D$

There's a clear first step here: Apply modus tollens to the two hypotheses to show that Alice can't swim. This means that they both can't swim, or, stated another way, everyone in our small domain cannot swim. That's not quite what we were asked to show, but an application of generalized De Morgan's laws will get us the rest of the way:

(1)	$S(\text{Alice}) \rightarrow S(\text{Bob})$	[Given]
(2)	$\neg S(\text{Bob})$	[Given]
(3)	$\neg S(\text{Alice})$	[1, 2, Modus Tollens]
(4)	$\neg S(\text{Alice}) \wedge \neg S(\text{Bob})$	[2, 3, Conjunction]
(5)	$\forall x(\neg S(x)), x \in D$	[4, Universal Generalization]
(6)	$\therefore \neg \exists x(S(x)), x \in D$	[5, Gen. De Morgan's Laws]

Line (4) isn't strictly necessary; we could justify the universally quantified statement directly from lines (2) and (3). But, as universal quantification means that the expression is true for every member of the set, putting both $\neg S(\text{Alice})$ and $\neg S(\text{Bob})$ on the same line – showing that every member of the domain makes $\neg S()$ true – allows the jump to the quantified expression be a little easier to see.

Guess what Alice and Bob's parents will sign them up to learn this summer (besides cryptography)?¹³

Getting tired of writing the same domain over and over? Example 69 demonstrates one way of speeding things up: Give the domain a single-letter

¹³'Alice' and 'Bob' are the traditional names used by cryptographers for the people exchanging encrypted messages. 'Eve' might soon be in the class, too ...

label. You can also write the domain once and state that it is used for the entire answer. We'll demonstrate this in the next example.

The first two examples in this section started with facts that, until this chapter, we would have written as propositions, and ended with conclusions that are quantified expressions. It's entirely possible to start an argument with a quantified expression and end with a thinly-disguised proposition.

Example 70:

Problem: The entire Sharma family is escaping the afternoon heat at the mall. Everyone at the mall has sore feet. Demonstrate that Diya Sharma has sore feet.

Solution: You know what to do first: Identify the predicates needed to express the hypotheses and the conclusion, and then express them.

(Note: The domain of x is "People" in this solution.)

Predicates: $S(x)$: x is in the Sharma family
 $M(x)$: x is at the mall
 $F(x)$: x has sore feet

Hypotheses: $\forall x (S(x) \rightarrow M(x))$
 $\forall x (M(x) \rightarrow F(x))$
 $S(\text{Diya})$

Conclusion: $F(\text{Diya})$

The shortest way to construct this argument is to start by combining the two quantified expressions. We can do this with hypothetical syllogism because they are both universally quantified and both have the same domain. Add in a little universal instantiation and a pinch of modus ponens, and we've got an argument. Because we've already stated that $x \in \text{People}$ for the entire answer, we won't include the domain in the argument.

(1)	$\forall x (S(x) \rightarrow M(x))$	[Given]
(2)	$\forall x (M(x) \rightarrow F(x))$	[Given]
(3)	$\forall x (S(x) \rightarrow F(x))$	[1, 2, Hypothetical Syllogism]
(4)	$S(\text{Diya})$	[Given]
(5)	$S(\text{Diya}) \rightarrow F(\text{Diya})$	[3, Universal Instantiation]
(6)	$\therefore F(\text{Diya})$	[4, 5, Modus Ponens]

For practice: Construct a second argument that also uses hypothetical syllogism, but later in the argument. Want more practice? Make a third argument that doesn't use hypothetical syllogism at all. Both of these versions will be longer than the argument given here, but validity also matters more than brevity.

This is a good time for a reminder: Don't overthink the context of your arguments, just as we warned you not to overthink expressions of predicates. Don't worry about details such as when the Sharmas are at the mall, or if Diya is even old enough to walk. Our purpose here is to make you comfortable with logical reasoning. Accept the context of the example, work within that framework, and you'll avoid introducing irrelevant details and concerns that will slow you down by complicating the problem unnecessarily.

3.5 Fallacies

Until now, our arguments have been firmly supported by logic. Detailing that support was a bit tedious, perhaps, but wasn't too difficult with the help of rules of inference and logical equivalences. Not many people have the knowledge or the patience to verify that their conclusions follow from a logical argument. An unfortunate consequence is that people can create – and worse, believe and act upon – invalid arguments that appear to be valid due to lack of critical examination. Such arguments are examples of *specious reasoning*.

specious reasoning

A noted creator of, and believer in, specious arguments is the animated character Homer Simpson in the long-running American television series “The Simpsons.” In the seventh season episode “Much Apu About Nothing,” in the scene pictured in Figure 3.3, Homer and his intelligent (and surprisingly well-educated) daughter Lisa have a conversation about the impact of a recently-implemented bear patrol in Springfield:




Figure 3.3: Lisa tries to explain specious reasoning to Homer in “Much Apu About Nothing,” a seventh season episode of “The Simpsons.” Credit: Fox.

Homer: Ah, not a bear in sight. The Bear Patrol must be working like a charm.

Lisa: That’s specious reasoning, Dad.

Homer: Thank you, honey.

Lisa: By your logic I could claim that this rock keeps tigers away.

Homer: Oh, how does it work?

Lisa: It doesn’t work.

Homer: Uh-huh.

Lisa: It’s just a stupid rock.

Homer: Uh-huh.

Lisa: But I don’t see any tigers around here, do you?

Homer: Lisa, I want to buy your rock.

Lisa is trying to demonstrate that correlation doesn’t mean causation. While it’s true (at least for that episode) that the rock is in their neighborhood and tigers are not, there is no support for the implication that the former implies the latter. When an argument is based upon an incorrect or

unsupported inference, the argument is a *fallacy*.

Definition 21: Logical Fallacy

A *logical fallacy* is an argument built with an incorrect inference.

logical fallacy

(Another name for logical fallacy is *propositional fallacy*.)

Be aware that there are also *informal fallacies*, those that demonstrate less specific errors of reasoning. We mention a few famous examples of such ‘illogical’ fallacies in section 3.5.2. First, let’s cover some logical fallacies.

informal fallacy

3.5.1 Logical Fallacies

Recall that all rules of inference are based on tautologies – when the hypotheses are true, the conclusion must also be true. Fallacies are often based on implications that mimic actual rules of inference, suggesting to the reader that the fallacy is also logically supported.

The following examples present three common logical fallacies.

Affirming the Conclusion

Example 71:

Problem: Does the following argument use a valid rule of inference?

Whenever it is sunny, Rodrigo wears sunglasses. Today he’s wearing a pair, therefore it must be sunny.

Solution: No, this argument doesn’t use a valid rule of inference – but the example does seem to make sense, because why else would anyone wear sunglasses? Actually, there are multiple reasons: The clouds are thin, it’s windy, Rodrigo has an unsightly stye, etc. Any of those can explain the wearing of sunglasses on a less-than-sunny day.

The fallacious argument of Example 71 can be expressed in logic notation like this:

$$\begin{array}{r}
 (1) \quad s \rightarrow g \\
 (2) \quad \underline{g} \\
 (3) \quad \therefore s
 \end{array}$$

*affirming the
conclusion*

In this form, it looks a lot like modus ponens, but of course it isn't. Because this fallacy claims the truth of the hypothesis (s) by assuming the truth of the conclusion (g), it is known by the name *Affirming the Conclusion* (a.k.a. *Affirming the Consequent*).

But why isn't this a valid argument? We are assuming that g is true, and that $s \rightarrow g \equiv s \rightarrow \mathbf{T}$ is also true. The problem is that the truth of $s \rightarrow \mathbf{T}$ doesn't mean that s must be true. In fact, $s \rightarrow \mathbf{T} \equiv \mathbf{T}$ regardless of the value of s ($\mathbf{T} \rightarrow \mathbf{T} \equiv \mathbf{F} \rightarrow \mathbf{T} \equiv \mathbf{T}$, by the definition of implication). We cannot conclude anything about the truth of s from the truth of those two hypotheses.

Another way to see that Affirming the Conclusion is a fallacy is by remembering that rules of inference are based on tautologies (see subsection 3.3.1). To show that $((s \rightarrow g) \wedge g) \rightarrow s$ is not a tautology, we need to find a circumstance in which conjunction of the hypotheses $((s \rightarrow g) \wedge g)$ is true but the conclusion (s) is false.

You could build a truth table to do this, but instead we'll just reason it through. For the logical AND to be true, both operands have to be true. That means g must be true and $s \rightarrow g (\equiv s \rightarrow \mathbf{T})$ must also be true, which it is when s is false. Thus, we've found a row in the truth table for which $((s \rightarrow g) \wedge g) \rightarrow s \equiv \mathbf{T} \rightarrow \mathbf{F} \equiv \mathbf{F}$, demonstrating that $((s \rightarrow g) \wedge g) \rightarrow s$ is not a tautology, and verifying that Affirming the Conclusion is a fallacy rather than a rule of inference.

Denying the Hypothesis

Remember those word analogy questions from standardized tests of vocabulary? They are often phrased like this: "This discrete math book is to fine literature as expired vinegar is to (a) overly-sweetened soft drinks, (b) discount floor wax, (c) the purest spring water, or (d) a bottle of 1787 Chateau Margaux?"¹⁴ That question type is applicable here. We just saw that Affirming the Conclusion is much like modus ponens. Similarly, the fallacy known as *Denying the Hypothesis* is much like modus tollens.

*denying the
hypothesis*

¹⁴A bottle of which, insured for \$225,000, was accidentally broken by its owner while he was showing it off to guests at a 1989 dinner in New York. Oopsie!

Example 72:

Problem: Does the following argument use a valid rule of inference?

You don't have a deck of cards. If you had one, you'd play poker.
So, no poker game for you.

Solution: Give yourself two demerits if you answered 'yes'. Sounds pretty good, though, doesn't it? Here's the argument in logical notation:

$$\begin{array}{l} (1) \quad \neg d \\ (2) \quad d \rightarrow p \\ \hline (3) \quad \therefore \neg p \end{array}$$

It's close to modus tollens, but not close enough. If you apply the Law of Contraposition to $d \rightarrow p$, you'll get $\neg p \rightarrow \neg d$, giving the argument exactly the same form as Affirming the Conclusion, and we've just shown that it isn't a valid argument.

Affirming a Disjunct

Our third logical fallacy, *Affirming a Disjunct*, is, as the name suggests, based on a disjunction rather than upon an implication. It confuses people because it is similar to disjunctive syllogism.

*affirming a
disjunct*

Example 73:

Problem: Does the following argument use a valid rule of inference?

To graduate with a computer science degree from Wossamotta U.,¹⁵ you need to successfully complete either Compilers or Operating Systems. You complete Compilers. Therefore, you did not complete Operating Systems.

Solution: Shockingly, the answer is still 'no.' Here's the form:

$$\begin{array}{l} (1) \quad c \vee o \\ (2) \quad c \\ \hline (3) \quad \therefore \neg o \end{array}$$

The problem with this one is in the implied assumption that you must choose exactly one of c or o to make $c \vee o$ true, when we know that there's a third way to make inclusive-OR true: $\mathbf{T} \vee \mathbf{T}$. That is, o could be true or false; we cannot claim that it must be false.

For thought: Can you turn this fallacy into a rule of inference by replacing the inclusive-OR with an exclusive-OR?

3.5.2 Informal Fallacies

informal fallacy

When an argument has an error of reason other than a poor logical inference, it is termed an *informal fallacy*. There are many of these; we offer two famous examples.

Begging the Question

You may have heard someone use a form of this phrase in conversation. For example: “The existence of seedless watermelons begs the question: How do they reproduce?” That’s a literal interpretation of the phrase “begs the question,” in which the speaker is imploring you to ask the suggested question.

As a fallacy, the phrase has a much different meaning: The argument’s conclusion is assumed true rather than shown to be true.

Example 74:

Problem: What’s wrong with the reasoning used in this argument?

Ladies and gentlemen of the jury: Let’s assume that Mr. Naidoo ate his mid-day meal of bunny chow¹⁶ on the day in question. Let us further assume that he paid in cash, as he claims, and left a generous tip. By his own admission, he was so sated that he didn’t eat again until the next morning. Ladies and gentlemen, you must agree: Mr. Naidoo ate that lunch!

Solution: You can almost believe that this was a portion of a lawyer’s concluding remarks to a jury, but even though it sounds good, you shouldn’t be convinced of the truth of the conclusion. The problem is that the jury

¹⁶Famous for having offered a football scholarship to B. J. Moose.

is asked to assume that the man ate lunch, and then to conclude that the man ate lunch. That is, the argument has this form: $\dots, p, \dots / \therefore p$.

Logically, this works – having assumed p to be true, we can correctly conclude that p is true under that assumption. But, practically, such an argument is a waste of our time.

Begging the Question is also known as *circular reasoning*, which makes sense given that we end up exactly where we started – our conclusion is one of our hypotheses.

circular reasoning

An interesting property of this argument form is that, like rules of inference, it *is* based on a tautology (specifically, $p \rightarrow p$). Thus, in terms of its structure, it's a valid argument, just not one that justifies any new truths. That's why Begging the Question is considered to be an informal rather than formal fallacy.

You may be wondering: Why is Simplification considered to be a rule of inference when it sure as heck looks like Begging the Question in a stylish Halloween costume? Yes, we could rewrite Simplification ($p \wedge q / \therefore p$) as $p, q / \therefore p$, because the notation p, q represents $p \wedge q$. However, the purpose of Simplification is to justify the truth of a single operand of the compound proposition $p \wedge q$ after we have discovered (or have been told) that $p \wedge q$ is true. If we are told (or have discovered ourselves) that p is true, we don't need any other justification for accepting p 's truth.

No True Scotsman

This informal fallacy can be viewed as a combination of Begging the Question and another, known as *equivocation*, in which the argument is based on an expression that has multiple meanings or interpretations. Because of these origins, the No True Scotsman fallacy isn't a distinct fallacy, but that doesn't lessen its popularity, or make it less worthy of consideration.

equivocation

Philosopher Anthony Flew is credited with the naming of this fallacy. His example involves sex,¹⁷ so to remain G-rated we will present a less-offensive version.

¹⁶A common South African dish. For humans. Ever have a stew served in a “bread bowl?” Same thing, really.

¹⁷Simmer down! It has the word ‘sex’ in it; that's as salacious as it gets.

Example 75:

Problem: What's wrong with the reasoning used in this argument?

Forbes was taken aback. Sitting in the local coffee house with his wife, Isla, he had just watched a neighbor pour sugar in his morning coffee. “No Scotsman puts *sugar* in his coffee,” Forbes muttered. Isla gave him a familiar look. “Forbes, love, your cousin Angus also takes his coffee with sugar.” Now Forbes was appalled. “I’m telling you, no **TRUE** Scotsman puts sugar in his coffee!”

Solution: To see this as a form of Begging the Question, notice that Forbes is assuming (in the face of mounting evidence to the contrary) that no Scotsman puts sugar in coffee, yet concludes the same thing, albeit with an insignificant additional word and an increase in volume. It's also equivocal, in that Forbes intends “Scotsman” and “true Scotsman” to have different meanings in an attempt to salvage his point.