LOGIC FOR BEGINNERS

"MEN ARE APT TO MISTAKE THE STRENGTH OF THEIR FEELING FOR THE STRENGTH OF THEIR ARGUMENT. THE HEATED MIND RESENTS THE CHILL TOUCH AND RELENTLESS SCRUTINY OF LOGIC."
—William E. Gladstone (1809-1898)

[4] ARGUMENT ANALYSIS

When most people hear the word 'argument', they think of disagreements or conflicts of some sort, as in:

1. "Bob and Alice are having a big argument about that new couch he bought yesterday."

Logic gives the word a rather different meaning: **An argument is a set of reasons that suggests the truth of some claim.** An argument is the laying out of the reasons (or evidence) for believing some claim. To analyze an argument is to attempt to understand those reasons to believe, and then to evaluate them as good or bad reasons.

It is in this latter sense of 'argument' that we say that **logic** is the study of good arguments or good reasoning, that is, logic is a study of how we *ought* to argue or reason, rather than of how we do in fact argue or reason. (In this regard, logic resembles ethics, which is a study of how we ought to behave, not how we in fact do behave.) Arguments are the primary objects of study in logic. A good argument will give reasons or support for believing some disputed or questioned claim. Bad arguments *seem* to give support when in fact they do not; these we call **fallacies**.

[Poem]

FIRE AND ICE

Some say the world will end in fire,
Some say in ice.
From what I've tasted of desire
I hold with those who favor fire.
But if it had to perish twice,
I think I know enough of hate
To say that for destruction ice
Is also great
And would suffice.

Robert Frost, 1920 (1874-1963)

PREMISES AND CONCLUSION

An **argument** is an attempt to support the truth of some statement (the **conclusion**), based on the alleged truth of one or more other statements (the **premises**). The premises, if they are to serve as premises, must be statements that everyone accepts. These premises then serve as reasons for accepting some additional claim — the **conclusion** — which is the statement in doubt, and thus in need of support. The most basic logical skill is being able to distinguish the premises from the conclusion. Consider the following passage:

> When individuals voluntarily abandon property, they forfeit any expectation of privacy in it that they might have had. Therefore, a warrantless search or seizure of abandoned property is not unreasonable under the Fourth Amendment. [Judge Stephanie Kulp Seymour, *United States v. Jones*]

The claim being argued for here is what follows the word 'therefore':

2. A warrantless search or seizure of abandoned property is not unreasonable under the Fourth Amendment.

That's the conclusion. And why should we believe this conclusion? The reason to believe the conclusion is provided in a single premise:

3. When individuals voluntarily abandon property, they forfeit any expectation of privacy in it that they might have had.

Normally you can discover the conclusion and premise(s) by a careful reading of the passage. What does the author want you to believe? (That will be the conclusion.) And what reasons does she give you to believe it? (Those will be the premises.) We often use certain indicator words to help highlight the reasoning involved (see the box listing common indicator words for premises and conclusions).

Here's another argument:

> Because both the world and a watch have multiple parts that seem to interact with one another for some purpose, and since watches were created by a designing intelligence, it follows that the world was probably also created by a designing intelligence.

It is generally helpful to mark up a passage, enclosing the various claims in parentheses and highlighting the indicator words:

> **Because** (both the world and a watch have multiple parts that seem to interact with one another for some purpose), and **since** (watches were created by a designing intelligence), **it follows that** (the world was probably also created by a designing intelligence).

Here, the first two claims are premises, and the conclusion that is presumably supported by these two premises is the last claim. Whether this conclusion actually follows from these premises will be considered more closely in a later chapter.

Conclusion-indicators
include:
therefore, wherefore, thus, consequently, hence, accordingly, entails that, for this reason, so, it follows that, as a result, suggests that, proves that, indicates that, is likely that.

Premise-indicators
include:
since, because, given that, in that, as indicated by, for, owing to, inasmuch as, may be inferred from, in view of the fact that.

SENTENCES AND THEIR USES

An argument consists of one or more premises and a conclusion; premises and conclusions are all statements; and a statement is a kind of sentence. So first we need to make sure we know what we mean by 'sentence' and 'statement', after which we can take a closer look at arguments.

A **sentence** is a series of words strung together following the rules of grammar. This is a sentence. Here's another one. And if we collect together all the sentences, from all the different languages, into one group, a sub-group of these sentences will be statements. A **statement** is a sentence with a **truth-value** (i.e., a sentence that is either true or false).

Sentences can be used to do various things, and in general they can perform any of three separate roles or *functions*:

- informative (e.g., "The world will end in two minutes.")
- expressive (e.g., "You can't be serious!")
- directive (e.g., "Please lift this piano off my nose.")

Only sentences meant informatively can be said to express statements, and so have a truth-value ("Ow!" or "Get out of here!" or "Where's my other sock?" do not have truth-values). In logic, we are normally concerned only with those sentences that have truth-values.

Apart from these functions, sentences come in various *forms* (this is the sentence's outward appearance), such as:

- *questions*: "Are you sitting on my hat?"
- *proposals*: "Let's buy the green one."
- *commands*: "Go to bed!"
- *exclamations*: "Nuts!"
- *assertions*: "There's coffee in the pot."

One might think that all assertions are informative, all exclamations expressive, and all commands and questions directive, but such a correspondence does not always hold. The following four sentences presumably all serve the same directive function of getting someone to bring more coffee, yet they all have different forms:

- Is there any coffee left? (question)
- Let's have another cup. (proposal)
- Bring me some coffee! (command)
- O cruel fate, that my cup is empty! (exclamation)
- I would like some more coffee. (assertion)

Finally, a sentence doesn't have to be true to be a statement. "The earth has five moons" is just as much a statement as "The earth has one moon" — although the first is false and the second true — and either might show up in an argument.

Statements and Propositions

Sentences with truth-values are statements. To be able to determine the truth-value of a statement, however, we need to understand its **meaning**, and we then need to check that meaning against reality, to see if the meaning corresponds with reality. If it does, then the sentence is true.

The meaning of a statement — when this meaning is unambiguous — is what we call the **proposition** (or **propositional-content**). Propositions are the **meaning** expressed unambiguously by statements, and (often ambiguously) by sentences. For instance, the sentence

1. His shirt is red.

has a truth value, and so is a statement, but we have no sense of the truth-value without first knowing which shirt is intended. Anyone who understands English will know that this statement is claiming that some male-gendered being has a shirt that is red; this is one level of meaning. Until we know who owns the shirt, however, and which shirt is being discussed, we won't be able to determine whether the statement corresponds with the way things really are, namely, that his shirt is in fact red.

In the following, I will use 'statement' in this fuller sense, as that set of words that unambiguously captures the propositional-content of a sentence, and which also has a truth-value. So, for instance, when I utter the sentence "I am sitting," this sentence — as a statement — would need to have the pronoun replaced with my name, and the time of the utterance would also have to be indicated, since sometimes I sit and sometimes I don't. In the following examples, we see the same statement expressed in three different sentences (1, 2, 3), and the same sentence expressing two different statements (1, 4)

	Sentence	**Statement**
1	**I am sitting.** [spoken by Steve Naragon at noon on June 3, 2016]	Steve Naragon is sitting at noon on June 3, 2016.
2	**He is sitting.** [spoken by another of Steve Naragon at noon on June 3, 2016]	Steve Naragon is sitting at noon on June 3, 2016.
3	**I was sitting at noon.** [spoken by Steve Naragon at 4pm on June 3, 2016]	Steve Naragon is sitting at noon on June 3, 2016.
4	**I am sitting.** [spoken by Alice Miller at noon on June 3, 2016]	Alice Miller is sitting at noon on June 3, 2016.

Different instances of the same sentence can have quite different meanings. For instance, the sentence "I am sitting" contains exactly three words, and it's the same sentence whether *I* say it or *you* say it. But the *meaning* of this sentence will change with the speaker, since the 'I' means Steve Naragon when I say it, but not when you say it. So here we have one sentence, but two statements, each expressing a different proposition. Finally, this sentence will sometimes be true, and sometimes false, depending on who says it and when they say it (namely, whether or not they are actually sitting when they say it). Until we know the full meaning, we won't be able to decide what statement is being expressed, and thus whether what is being said is true.

One important role of language is to convey some propositional meaning from one person (the sender) to another (the receiver). Depending on the beliefs and backgrounds of these two parties, the intended meaning of the sender can differ considerably from the meaning understood by the receiver. This is true in any number of senses, but also in one very direct sense: When someone says, for instance, that "Heather's dog has rabies," any competent English-speaker will understand what is meant, since they will know the meaning of each word of the sentence; but they won't have the fullest understanding of this sentence — and thus the proposition being expressed — until they know exactly which dog is claimed to have rabies. The sender and receiver might have in mind quite different Heather's, or perhaps the same Heather but different dogs (should more than one dog live with Heather).

Simple and Complex Statements

The statements comprising the premises and conclusion of an argument can be either simple or complex. A **simple statement** cannot be analyzed into simpler statements, while a **complex statement** can. This is not as clear-cut as you might imagine, but for our purposes we shall assume that it *is* clear-cut. So, for instance,

1. A hedgehog is climbing up the back of your trousers.

will, for our purposes, count as a simple statement, even though a pickier person might conceivably analyze it into something like this:

2. There is at least one hedgehog.
3. There is at least one person (namely, you).
4. There is at least one pair of trousers.
5. You are standing in a certain relationship to those trousers (namely, you are wearing them).
6. The hedgehog is standing in a certain relationship to the trousers as described in statement 5 (namely, it is climbing up the back of them).

I will not expect such a high level of pickiness here.

> **Liberal Education**
>
> "Liberal education in our colleges and universities is and should be Socratic, committed to the activation of each student's independent mind and to the production of a community that can genuinely reason together about a problem, not simply trade claims and counterclaims."
>
> — Martha Nussbaum, *Cultivating Humanity* (1997), p. 19.

A **complex statement** is any statement that *can* be analyzed into simpler statements. Three common types of complex statements are disjunctive (p or q), conjunctive (p and q), and conditional statements (if p, then q). Let's first consider a **disjunctive statement**, which is marked by the terms "either … or …":

1. Either that's a hedgehog climbing up your trousers or it's a porcupine.

can be analyzed into the two simple statements:

2. A hedgehog is climbing up your trousers.
3. A porcupine is climbing up your trousers.

and we could re-write the original sentence in the following form, which is more straight-forward but less elegant, enclosing each simple statement in parentheses:

4. (A hedgehog is climbing up your trousers) **or** (A porcupine is climbing up your trousers.).

Sometimes complex statements will appear to be simple statements, like the following **conjunctive statement**:

5. Ed and Bob are Republicans.

On closer inspection, it is found to be analyzable into two simple statements:

6. Ed is a Republican.
7. Bob is a Republican.

... and can be more perspicuously written as:

8. (Ed is a Republican) **and** (Bob is a Republican).

Conjunctive statements usually include a conjunction like ‘and’ or ‘but’.

Finally, some of the most common arguments that we encounter include inferences drawn from a **conditional** statement, like the following:

9. If you mow my lawn this afternoon, then I’ll pay you $20.

This can be re-written with the simple statements enclosed in parentheses, in order to make the form more apparent:

10. **If** (you mow my lawn this afternoon), **then** (I’ll pay you $20).

In general, every conditional statement has the form “If p, then q” — where ‘p’ and ‘q’ are statements. Similarly, every disjunctive statement has the form “p or q” and every conjunctive statement has the form “p and q.”

[extra]

Truth Functionality

One important feature of these complex statements is that their truth-value is dependent upon the truth-values of the simple statements that comprise them. So for any disjunctive statement, ‘p or q’, the statement will be true so long as one or both of the simple statements is true; if both p and q are false, then ‘p or q’ is false. Any conjunctive statement, ‘p and q’, is true so long as both p and q are true; otherwise it is false. And any conditional statement will be true except for one situation: if p is true and q is false, then the conditional statement ‘if p, then q’ is false; otherwise it is true. It is for this reason that we call these complex statements **truth functional**: their truth-value is a function of the truth-values of their parts.

Necessary and Sufficient Conditions

Necessary and sufficient conditions are often explained in terms of conditional statements, and so it is useful to discuss them here, even though the fit is not as tidy as is sometimes presumed. On the surface at least — and I plan to go no deeper than that here — these conditions are fairly straight-forward, and can be captured with a simple conditional statement, ‘**if p, then q**’: p is a **sufficient condition** of q, and q is a **necessary condition** of p. This suggests a symmetry between the two conditions, viewing them as essentially the same relationship, but moving in opposite directions. So, for instance,

1. If something is a square, then it has four sides.

Being a square is a sufficient condition for having four sides, and having four sides is a necessary condition for being a square. Having four sides is also a necessary condition for being a rectangle (or any other quadrilateral), and being a square is also a sufficient condition for being a two-dimensional figure. Having a daughter is a *sufficient condition* for being a parent, but not a necessary condition (since you might instead have a son).

For any given object or event, we can speak of its various necessary and sufficient conditions. A necessary condition (NC) of some object or event is whatever *has* to happen, given the object or event, while a sufficient condition (SC) of some object or event is adequate, all by itself, to bring about that object or event.[5]

X	**... is a ...**	**... of ... Y**
Having four sides	NC (but not SC)	Being a square
Being a square	SC (but not NC)	Having four sides
Having a daughter	SC (but not NC)	Being a parent
Being a parent	NC (but not SC)	Having a daughter
Being an unmarried man	NC and SC	Being a bachelor
Wearing blue suede shoes	neither NC nor SC	Being a college professor.

This all seems fairly straight-forward, at least until I show up at the cinema wanting to watch a movie, and the person at the ticket counter informs me that:

2. If you want to see this movie, then you need to buy a ticket.

[5] To put this in terms of statements: A is a sufficient condition of B if A’s truth guarantees B’s truth, and B is a necessary condition of A if B must be true whenever A is true.

This seems to mean that buying a ticket is a necessary condition for seeing the movie (assuming that I remain legal, and don't somehow sneak into the theater without one); but can we plausibly say that either is a condition of the other? I can want to see a movie and still not buy a ticket, and my wanting to see a movie is not sufficient for buying the ticket (I also need the money). The standard English used here does not match the logical nature of the 'if…then…' construction, and would need to be written as:

3. If you see this movie, then you will buy (or 'will have bought') a ticket,

which doesn't quite capture the sense of (2).

Or suppose I decide that …

4. If Tom gets an A in Calculus, then we'll celebrate.

Here it seems clear that Tom's A is a sufficient condition for our celebration; but it sounds odd, if not downright false, to claim as well that our celebration is a necessary condition for Tom's A in the class.

I leave these as unresolved puzzles, hanging from the deep complexity of the English word 'if' (and the concept of cause that often lurks beneath).

Arguments vs Explanations

Arguments and explanations typically share the same structure, but they differ in the direction of the reasoning. With **arguments**, the conclusion is always the statement that is known less well, and it is supported by appealing to statements that are better known (the premises). With **explanations**, the thing or event being explained (the *explanandum*) is best known, and the explanation of that thing or event — why or how it happened (the *explanans*) is less known. Consider the following sentences:

1. Mary broke the window because she wasn't paying attention while pitching the ball.
2. Mary broke the window because she forgot her key and needed to get in.
3. Mary broke the window because she was the only person nearby at the time.

The first two sentences are explanations, while the last is an argument (statements in bold-face are what we appear to know better):

Mary wasn't paying attention.	Mary forgot her key.	**Mary was the only person nearby.**
Mary broke the window.	**Mary broke the window.**	Mary broke the window.

In all three of these columns, the first sentence could be introduced with a 'because' or a 'since,' etc. But in the first two columns, it is the second sentence that is known best, while the first sentence is offered as a conjecture. In the first two sentences, the "because"-phrase explains why Mary broke the window; in the third sentence, the "because"-phrase gives us a reason to believe that Mary broke the window.

When the *explanans* is controversial or uncertain, then it is arrived at through a kind of reasoning called hypothetical induction, which will be discussed below.

Arguments vs Expository Passages

Passages containing arguments and explanations are also to be distinguished from merely **expository passages**: these will describe, develop, or illustrate some point, rather than offer reasons to believe it. Consider the following two paragraphs:

> Musical composition is one of the most mysterious of all art forms. People who can easily come to terms with a work of literature or a painting are still often baffled by the process by which a piece of music — appearing in material form as notation — must then be translated back into sound through the medium of a third party — the performer. Unlike a painting, a musical composition cannot be owned (except by its creator); and although a score may be published, like a book, it may remain incomprehensible to the general public until it is performed. Although a piece may be played thousands of

> times, each repetition is entirely individual, and interpretations by different players may vary widely. [Wade-Matthews and Thompson, *The Encyclopedia of Music*]

> There are three familiar states of matter: solid, liquid, and gas. Solid objects ordinarily maintain their shape and volume regardless of their location. A liquid occupies a definite volume, but assumes the shape of the occupied portion of its container. A gas maintains neither shape nor volume. It expands to fill completely whatever container it is in. [Hill and Kolb, *Chemistry for Changing Times*, 7th ed.]

Both of these paragraphs have a topic sentence (viz., the first sentence of each paragraph) that introduces the topic for the paragraph. In the first, the sentences following the topic sentence give the reader reasons to believe that "musical composition is one of the most mysterious of all art forms." Because of this, the first paragraph is readily understood as containing an argument. In the second paragraph, the topic sentence is not supported by the remainder of the paragraph (none of the sentences give a reason to believe that "there are three familiar states of matter"). Instead, they are merely developing the topic sentence by describing these three states.

ENTHYMEMES

An enthymeme is an argument that is missing either a premise or a conclusion, or both. In everyday life, most arguments are presented as enthymemes, because one or more of the premises is so obvious that explicitly mentioning it would seem redundant, or else because the inference from the stated premises to the conclusion is so obvious that stating it would insult the intelligence of the reader. For example …

1. We cannot trust this man, for he has lied in the past.

Here we find an argument with a conclusion [C] and one premise [Pr] ...

2. [C] (We cannot trust this man), for [Pr1] (he has lied in the past).

... but it seems to be missing a second premise necessary for the conclusion to follow, namely,

3. [Pr2] Whoever has lied in the past cannot be trusted.

When filling out these enthymemes, we need always to follow the **principle of charity**, which advises us to assume that the person making the argument is rational and not arguing invalidly; thus, if at all possible, the missing statements need to be inserted in such a way that the resulting argument is as good as possible.

Enthymemes are generally a good thing. Spelling out every premise, or always explicitly stating the conclusion, can solidify an otherwise fluid prose into a nearly immovable slab of text. In a witty poke at the operatic composer Richard Wagner (1813-1883), Mark Twain wrote:

> There is no law against composing music when one has no ideas whatsoever. The music of Wagner, therefore, is perfectly legal.

Twain could have expressed himself more perspicuously with the following …

(1) It is not illegal to compose music barren of any ideas.	[stated premise]
(2) Richard Wagner has composed music barren of any ideas.	[assumed premise]
(3) Therefore, Richard Wagner has done nothing illegal in composing this music.	[stated conclusion]

… but it is to our better fortune that he did not, nor would it have been as humorous.

The above is an honorable form of enthymeme, but some are dishonorable, for they are sometimes used to cloak dubious premises. This is common in advertising, where loose associations are often made between, say, using a certain mouthwash and enjoying a happy life filled with beautiful and adoring friends. The implied premise necessary to get to the desired conclusion would look something like this: "If you use our mouthwash, then your life will soon be filled with beautiful and adoring friends." When stated so openly, however, the argument loses all ability to persuade.

[5] EXTENDED ARGUMENTS

The deductive and inductive arguments discussed below will all be simple arguments — a set of statements, one of which is the conclusion, whose truth is supported by the remaining statements that are the premises, or reasons supporting the truth of the conclusion. Most arguments that we encounter in everyday life, however, are rather more complicated than this, typically involving arguments nested inside other arguments, and with multiple lines of support. These more complex arguments we call **extended arguments**.

Our basic task, with these **extended arguments**, is the same as with simple arguments: (1) identify the conclusion (the statement that the arguer is wanting us to believe, and towards which all the other statements are trying to point) and (2) evaluate the extent to which the premises support the truth of the conclusion. The new problem here is deciding how to map out the more complicated structure of logical support. For this we need a few more terms. There are four basic patterns to be found in any extended argument, and these can themselves be mixed and matched: horizontal, vertical, conjoint premises, and multiple conclusions. The last pattern (multiple conclusions) is less common.

The **horizontal pattern** is where two or more premises give independent support to a conclusion (that is, each premise, all by itself, supports the conclusion). The **vertical pattern** is where one statement supports a second statement, the second statement supports a third, and so on. The **conjoint premises pattern** is where two or more premises support the conclusion, but only when grouped together. Every deductive and inductive argument mentioned below will have this pattern. Finally, the **multiple conclusions pattern** is where, at the end of the argument, one finds more than one final conclusion being supported. (Note that most extended arguments will have many conclusions, where a conclusion is then used as a premise for another conclusion; but only some will have multiple *final* conclusions.)

When tracing out the support structure of an extended argument, keep asking the question: "Why should I believe this?" Number the statements in the passage, then look at each statement and ask: "What other statement in the passage gives me reason to believe **this** statement?" If any of the other statements speak to this question, then start a diagram on a scrap of paper and, using numbers and arrows, draw an arrow from the supporting premise to the conclusion. It may take several run-throughs, but eventually the support structure will become clear, and in the end you will have a much better understanding of the argument and whether it really does offer adequate grounds for believing the final conclusion. What follows are extremely abbreviated arguments to serve as examples of these four patterns.

Horizontal

> (1) We should not build more nuclear power plants in the United States. (2) Nuclear power is a dangerous technology to those presently living, (3) it places an unfair burden on future generations, and (4) we don't really need the additional power such plants would generate.

In this passage, (1) is supported separately by (2), (3), and (4).

Vertical

> (1) Nuclear power is a dangerous technology. (2) Accidents at these power plants are inevitable, since (3) human error will sooner or later lead to a mistake. (4) We should not build more nuclear power plants in the United States.

In this passage, (4) is supported by (1), which is supported by (2), which is supported by (3).

Conjoint

> (1) Nuclear power is an unacceptably dangerous technology. (2) Accidents at these power plants are inevitable and (3) accidents of this sort generally involve catastrophic consequences.

In this passage, (1) is supported conjointly by (2) and (3). In other words, neither (2) nor (3) offer any support to (1) except when they are combined together — for if the accidents are inevitable but trivial, then that does not suggest that nuclear power is a dangerous technology. Similarly, if the accidents are catastrophic, but avoidable, then nuclear technology need not be considered dangerous. But when these are combined, then the conclusion is supported.

Multiple Conclusions

> (1) Nuclear power is a dangerous way to generate electricity. Therefore, we should (2) stop building such plants, (3) support the development of solar technologies, and (4) do a better job conserving energy use.

In this passage, statements (2), (3), and (4) are all supported by (1).

[6] DEDUCTIVE LOGIC

All arguments that we come across in our day-to-day lives are either **deductive** or **inductive**. Good *deductive* arguments guarantee the truth of the conclusion, so long as the premises are true. Good *inductive* arguments, on the other hand, suggest that the conclusion is more likely true than false (and therefore more reasonable to believe than to disbelieve) if the premises are true. In what follows, we will first consider some deductive arguments, and then look at four kinds of inductive argument.

SOUNDNESS, VALIDITY, AND TRUTH

A good deductive argument is called a sound argument, and a **sound argument** is defined as one that (a) is *valid*, and (b) has only *true premises*.

A deductive argument is **valid** if the conclusion necessarily follows from the premises, in other words, if it is impossible for the premises to be true and the conclusion false. (This is a technical definition that you need to both memorize and understand!) Validity (or its opposite: invalidity) is a property of arguments, *not* of individual statements. Validity concerns the form or structure of the argument, rather than its content; specifically, validity concerns the inference-relation between the premises and the conclusion. We can learn whether an argument is valid or invalid simply by looking at the form — the way the parts are connected together — without knowing what the actual parts are at all. A good structure is valid and a bad structure is invalid. Validity is that property of an argument's structure that guarantees the truth of the conclusion given the truth of the premises. Another way of thinking about validity is as a *truth-preserving* property: if an argument is valid, then any truth contained in the premises will be preserved in the conclusion.

Soundness and unsoundness are also properties of arguments (and only of arguments), because they include validity as one of their components. Truth and falsity, on the other hand, are properties of individual statements (for instance, properties of the premises and the conclusion). Each statement has a *truth-value*, being either true or false.

EVALUATING ARGUMENTS

As we've seen above, a good deductive argument has two components: validity and true premises. Consequently, in evaluating a deductive argument, you need to ask two things: "Is the argument valid?" and "Are all of the premises true?" If the argument fails with either of these, then it fails completely. This should be your strategy whenever you are confronted by an argument: check its validity and check the truth of each premise.

How do you know when a premise is true or false? Sometimes this is obvious: "The earth consists entirely of discarded chewing gum" is false based on simple perceptual evidence (namely, the earth doesn't *appear* to consist of discarded chewing gum). Sometimes, however, we need to argue for the truth or falsity of a premise; in other words, deciding the truth of some premise might require you to evaluate a prior argument that has this premise as its conclusion. (From this you can see now how evaluating arguments can quickly grow quite complicated.) Often you won't be able to determine whether a premise is *certainly* true or *certainly* false; at best you will be able to say whether there are better reasons for believing a premise than for disbelieving it. All of this affects the strength of the argument as a whole.

Evaluating the validity of an argument can be a bit trickier, although people routinely do this informally (although they are often unaware of what exactly they are criticizing). The informal method for checking validity is called "**the method of counterexample**," whereby you imagine some scenario in which an argument with the *same form* as the argument in question has *obviously true premises* but an *obviously false conclusion*. Since validity guarantees a *true* conclusion from true premises, such an argument form is immediately shown to be invalid.[6]

Examples of evaluating the validity of arguments

Consider the following argument: "If it were raining, then you would need your umbrella; but it's not raining, so you don't need your umbrella." This is a simple argument with two premises and a conclusion. We could re-write it as follows:

(1) If it's raining, then you need your umbrella.
(2) It's not raining.
(3) ∴ You don't need your umbrella. ['∴' is a sign used to introduce a conclusion; read it as 'therefore']

In evaluating this argument's soundness, we need to investigate its validity and the truth of the two premises. Here the truth is relatively easy to determine, but what about the validity? There are two kinds of counterexample that you might offer. First, try to imagine some **alternate scenario** in which the premises of the argument will be true, but the conclusion false. A little reflection shows that we can do this with the above argument. For instance, it could well be true that you need your umbrella when it's raining (assuming that you don't want to get wet), and it could well be true that it isn't raining; but you might nonetheless need your umbrella, for perhaps it is snowing and you want to keep the snow off your head, or you need an umbrella to fend off a vicious dog or a mugger. Merely that it's *possible* for the premises to be true and the conclusion false proves the invalidity of the argument.

An improved form of the method of counterexample is where we strip away the content of the premises and conclusion, leaving only the form of the argument, and then find **substitution instances** that will make all the premises true but the conclusion false. Recall that validity concerns only the form of the argument, and not its content, so stripping away the content actually makes the evaluation of validity more straightforward. The form of the above argument is found in the recurring phrases. If we let 'p' stand for 'It's raining' and 'q' stand for 'You need your umbrella', then we can write the form of this argument as follows:

(1) If p, then q	Key:
(2) not-p	p = 'it's raining'
(3) ∴ not-q	q = 'you need your umbrella'

[denying the antecedent / **invalid**]

Let's now see if we can find substitution instances for 'p' and 'q' that will make the premises obviously true and conclusion obviously false. Substitute "That's a dog" for 'p' and "That's a mammal" for 'q', and now choose some cow as the subject of the argument: It's clearly true that if something is a dog, then it will be a mammal; and it's clearly true that the cow being discussed is not a dog; but it's also clearly false that the cow is not a mammal. The fact that we are able to find substitutions for 'p' and 'q' making the premises clearly true and the conclusion clearly false proves that the argument is invalid (indeed, any argument with that form will be invalid).

This particular argument form (if p then q; not-p; therefore, not-q) is such a common fallacy that it has even been given a name: *Denying the Antecedent*. This refers to the conditional ('if…then') statement: the 'p' is the antecedent

[6] This method can only prove invalidity. It cannot prove that an argument is valid, since the failure to discover a counterexample might be the result of your poor imagination rather than of the argument's structure.

and the ‘q’ is the consequent of the statement. In this fallacy the antecedent is denied, from which is illicitly inferred the denial of the consequent.

Now consider this argument: “If God exists, then there is no evil in the world, and God *does* exist; therefore, there is no evil in the world.” Written in standard form, the argument will look like this:

(1) If God exists, then there is no evil in the world.	If p, then q
(2) God exists.	p
(3) ∴ There is no evil in the world.	∴ q

[modus ponens / **valid**]

This is a valid argument, and goes by the name of ‘modus ponens’. Because the method of counterexample cannot prove validity, you’ll have to trust me on this; but if you want to idle away an afternoon, try to find a counterexample or substitution instance that will prove this form to be invalid. (Good luck.) Whether this argument is sound, however, is another matter, since neither of the premises is obviously true (at least not true without further argumentation). Many theists will reject (1), all atheists will reject (2), and most people reject the conclusion.

Now consider the following argument and its form:

(1) If the medicine doesn’t work, then the patient will die.	If p, then q
(2) The patient died.	q
(3) ∴ The medicine did not work.	∴ p

[affirming the consequent / **invalid**]

Let ‘p’ and ‘q’ stand for the same substitutions as above (“That animal is a dog” and “That animal is a mammal”), and once more consider a cow. As with denying the antecedent, both premises turn out to be true and the conclusion false, and so the argument is invalid. This fallacy is called *affirming the consequent*.

The French philosopher and scientist **René Descartes** (1596-1650) once wrote that “I cannot be identical with my body, since if I were, both my existence and my body’s existence would be equally dubitable, and they are not” — for he can doubt whether his body exists, but not that he (as a thinking thing) exists. Re-written, the argument looks like this:

(1) If I am identical with my body, then my existence and my body’s existence are equally dubitable.
(2) My existence and my body’s existence are not equally dubitable.
(3) ∴ I am not identical with my body.

This argument has a common — and valid — form called *modus tollens*:

(1) If p, then q
(2) not-q
(3) ∴ not-p

[modus tollens / **valid**]

Is it a good, or sound, argument? That depends on whether the premises are true. Descartes does a good job of demonstrating the truth of (2) in the first and second meditations of his *Meditations on First Philosophy* (1641), but the first premise is very likely false. After all, I might feel quite certain that Mark Twain wrote *The Adventures of Huckleberry Finn*, but feel less certain that Samuel Clemens did (this would be true for anyone who didn’t know that Mark Twain is the same man as Samuel Clemens). My *belief states* being different doesn’t require that *the objects of belief* differ. So Descartes’ argument is unsound; and it is unsound even if the conclusion turns out to be true. Remember: soundness is strictly a property of arguments, and while sound arguments always have true conclusions, unsound arguments can have either true or false conclusions.

Take a moment to study the following three arguments, all of which are valid (the first two have the form *modus ponens* and the third *modus tollens*). The first argument is unsound because its first premise is false.

VALID (modus ponens), but UNSOUND

(1) If you have dark hair, then you are Italian.	If p, then q

(2) You have dark hair. p
(3) ∴ You are Italian. ∴ q

VALID (modus ponens), and SOUND
(1) If you are human, then you are mammalian. If p, then q
(2) You are human. p
(3) ∴ You are mammalian. ∴ q

VALID (modus tollens), and SOUND
(1) If the medicine doesn't work, then the patient will die. If p, then q
(2) The patient did not die. not-q
(3) ∴ The medicine worked. ∴ not-p

As noted above, validity is a property of arguments that preserves the *truth* of premises when moving from premises to conclusion but, surprisingly, it does not preserve any of the *falseness*. If an argument is valid *and* if all the premises are true, then the conclusion is guaranteed to be true; but if one or more (or all) of the premises are false, the conclusion could be either true or false. Validity preserves only truth.

Proving Validity with Truth Tables

The proof of the validity of argument forms in propositional logic had to wait nearly 2000 years for the independent work of Emil Post (1897-1954) and Ludwig Wittgenstein (1889-1951). Wittgenstein developed, in his classic early work, the *Tractatus Logico-Philosophicus* (1921), an account of truth that is still in use today. A truth table, in effect, makes a complete list of every possible combination of truth-values of the component statements, and from that allows us to determine whether an argument is valid (or "truth-preserving") by seeing whether there is any possible combination in which all the premises are true but the conclusion false.

This might sound complicated, but it is actually quite simple and intuitive once you reflect on an example or two. Consider the argument: "I told you that if I won the lottery, I'd buy you dinner, and I did win the lottery; so I'm going to buy you dinner." Because this is such a straightforward argument, we can *see* that it's valid; but it's nice to be able to *demonstrate* its validity as well, especially since there are plenty of more complicated arguments where the validity is anything but straightforward, and where some mechanical method for demonstrating validity is crucial. This argument is an example of modus ponens:

(1) If I win the lottery, then I'll buy you dinner. If p, then q
(2) I won the lottery. p
(3) ∴ I'll buy you dinner. ∴ q

In this argument, there are two statements, p and q, and each of these might be either true or false. That means there are four possibilities — or possible worlds (ways the world might have turned out) — that we need to consider:

(1) I win the lottery and buy you dinner
(2) I win the lottery and don't buy you dinner.
(3) I don't win the lottery but still buy you dinner.
(4) I don't win the lottery and I don't buy you dinner.

We can summarize these possibilities in the first truth table (see "Truth Table - MP"), with each line of the truth table representing one of the four possible worlds.

This truth table shows the four possible combinations of the truth-values for p and q (the first two columns), and then lists the premises and the conclusion of the argument given above (the last three columns). When constructing a truth table to check the validity of an argument, you first fill in the possible combinations for the variables 'p' and 'q', and then you assign the truth-values for the premises and the conclusion for each of the four possible worlds. With modus ponens arguments, the second premise is just 'p', and so the truth-values will be the same as column one;

Truth Table - MP

	p	**q**	**p ⊃ q**	**p**	**q**
(1)	T	T	T	T	T
(2)	T	F	F	T	F
(3)	F	T	T	F	T
(4)	F	F	T	F	F
			P1	**P2**	**C**

likewise with the conclusion, which is identical to the 'q' column. The only column that might be puzzling is how to decide when a conditional statement is true. As it turns out, conditional statements ('if p, then q') are **false** only when the p-statement is true and the q-statement is false (this is the situation we find on row 2 of *Truth Table - MP*); otherwise they are true. For instance, the conditional statement "If I win the lottery, then I'll buy you dinner" is clearly false if I do in fact win the lottery and then refuse to buy you dinner. And it's clearly **true** if I both win the lottery **and** I buy you dinner; but it might seem a little odd that the statement is true even in those situations where I don't win the lottery. But think about it: if I buy you dinner even when I don't win the lottery, does that make the conditional statement false? Not at all! It just means I'm a really nice guy, and that I have money to spare, even without winning the lottery.

Truth Table - AC

	p	**q**	**p ⊃ q**	**q**	**p**
(1)	T	T	T	T	T
(2)	T	F	F	F	T
(3)	F	T	*T*	*T*	*F*
(4)	F	F	T	F	F
			P1	**P2**	**C**

After we have filled in the truth-values for the premises and conclusions, we are then ready to use the completed truth table to decide whether the argument is valid. Remember that an argument is valid only if it's impossible for the premises to all be true and the conclusion to be false. So to check validity, we need to examine each row in which all the premises are true; if any of these rows have a false conclusion, then we know that the argument is invalid, otherwise it's valid. If the conclusion is always true on those rows in which all the premises are true, then the argument is valid. With modus ponens, we see that the conclusion is true on every row where all of the premises are true (there is only one row where this is the case, namely row 1)

This is not the case with the argument form known as "affirming the consequent" (see *Truth Table - AC*). Here we see that both premises are true on rows 1 and 3, and that while the conclusion is true on row 1, it is false on row 3, proving that it is possible for the premises to all be true, yet the conclusion be false. This proves that any argument with the form of affirming the consequent is invalid.

A Few Standard Forms for Deductive Arguments

Examples of a few of the more common argument forms are given below. Use the following list of forms to parse these arguments into their component p's and q's, and then check your work with the answers in the footnotes below.

Modus Ponens (valid)
(1) p ⊃ q
(2) p
(3) ∴ q

Modus Tollens (valid)
(1) p ⊃ q
(2) ~q
(3) ∴ ~p

Hypothetical Syllogism (valid)
(1) p ⊃ q
(2) q ⊃ r
(3) ∴ p ⊃ r

Disjunctive Syllogism (valid)
(1) p v q
(2) ~p
(3) ∴ q

Denying the Antecedent (invalid)
(1) p ⊃ q
(2) ~p
(3) ∴ ~q

Affirming the Consequent (invalid)
(1) p ⊃ q
(2) q
(3) ∴ p

p ⊃ q reads 'if p, then q'
p v q reads 'p or q' (inclusive)
p · q reads 'p and q'
~p reads 'not p'
∴ reads 'therefore'

Modus Ponens (MP)

"If you have dark hair, then you are Italian. John has dark hair. Therefore, John is Italian."
[If p, then q; p; therefore, q.]

Modus Tollens (MT)

"If God exists, then there is no evil in the world. But there *is* evil. So God does not exist."
[If p, then q; not q; therefore, not-p.]

Hypothetical Syllogism (HS)

> "If human beings consist entirely of matter, then all of their actions are determined. If all of their actions are determined, then none of their actions are free. Therefore, if human beings consists entirely of matter, then none of their actions are free."
> [If p, then q; if q, then r; therefore, if p, then r.]

Compare with the following invalid argument:

> "If you're a brain surgeon, then you're wealthy. And if you're a brain surgeon, then you've also spent years and years in school. Therefore, if you're wealthy, then you've spent years and years in school."
> [If p, then q; if p, then r; therefore, if q, then r.]

Disjunctive Syllogism (DS)

> "The mind is either identical to the brain or else it is an immaterial substance. It is not identical to the brain. Therefore, it is an immaterial substance."
> [p or q; not-p; therefore, q.]

This is a valid argument, although it may not be sound, since the first premise is probably false — there is probably some third alternative — and of course the second premise would need to be supported with an additional argument. Compare with the following *invalid* argument:

> "This Intro class is going to be either really difficult or really interesting. It's turning out to be really interesting, so I guess it won't be very difficult."
> [p or q; q; therefore, not-p.]

What is happening in the second argument is that one of the alternatives is believed to be true, and from that it is wrongly inferred that the other alternative must be false — but the first premise claims only that one alternative is true, not that only one is true (for they might *both* be true).

Denying the Antecedent (DA; *invalid*)

> "If this is not the best of all possible worlds, then God does not exist. But this *is* the best of all possible worlds. So God *does* exist."
> [If p, then q; not-p; therefore, not-q.]

Affirming the Consequent (AC; *invalid*)

> "If the medicine doesn't work, then the patient will die. The patient died. So the medicine did not work."
> [If p, then q; q; therefore, not-p.]

Mixed (DS, MT, HS)

> "The world is either finite or infinite in space. If it's finite, then we would be able to arrive at its edge. But we cannot even conceive of doing this (much less doing it), so it is not finite. Therefore, it must be infinite."
> [p or not-p; if p, then q; not-q; therefore, not-p.]

> "The world is either finite or infinite in age. If it's infinite, then it had no beginning in time; and if the world has no beginning in time, then an infinite series has been completed. But an infinite series cannot be completed (by definition). Therefore, the world is not infinite; and so therefore the world's age is finite (i.e., it has a beginning)."
> [p or not-p; if not-p, then q; if q, then r; not-r; therefore, not-not-p; therefore, p.]

> "The world is either finite or infinite in age. If it's finite, then it had a beginning in time; and if the world had a beginning in time, then it was preceded by some empty time. But nothing can begin from an empty time. Therefore, the world is not finite; and so therefore the world's age is infinite (i.e., it has no beginning)."
> [p or not-p; if p, then q; if q, then r; not-r; therefore, not-p.]

Dilemma

Dilemmas can be analyzed into simpler argument forms (they make use of a combination of disjunctive syllogism and modus ponens or modus tollens), but they are such a common and effective rhetorical move that they are commonly treated as a distinct form of logical argument. Dilemmas are traditionally thought to have two "horns" (the two conditional statements), such that the opponent can be "impaled on the horns of a dilemma." Think of the dilemma as though it were a charging and angry bull. Being impaled on the horns of the dilemma is something to be avoided, whenever possible, and this is done either by "grasping one of the horns" or by "going between the horns." The first strategy is to question the truth of one of the conditional statements; the second is to contest the truth of the disjunctive claim (normally by identifying some third possibility that is more plausible than either of the given alternatives). Finally, dilemmas are sometimes counteracted by posing a counter-dilemma. This last maneuver is rhetorically effective, primarily because it tends to dazzle the opponents, although there's little reason to believe that doing this moves one any closer to the truth. Consider this dilemma:

> "If God cannot prevent evil, then he is not omnipotent. If God does not want to prevent evil, then he is not perfectly good. God either cannot or does not want to prevent evil (since evil exists). Therefore, God is either not omnipotent or not perfectly good."
> [If p, then q; if r, then s; p or r; therefore, q or s.]

A common response to this argument is to attack the right horn, and argue that there could well be reasons why God might allow a creature to suffer that are compatible with God's goodness. One might also attempt to go between the horns, by arguing that evil does not exist (both of these strategies will be explored in a later chapter).

In the next dilemma, going between the horns is not possible:

> "If test animals are similar enough to humans to make the research relevant, then experimenting on them is morally similar to experimenting on humans, and thus wrong. If test animals are too dissimilar, then the research is irrelevant. Either test animals are similar enough to humans or not. So experimenting on animals is either morally wrong or irrelevant."
> [If p, then q; if r, then s; p or not-p; therefore, q or s.]

Going between the horns is not possible, because the disjunctive claim is a necessary truth (p or not-p). The best approach is probably to contest the left horn, focusing on the different ways of being *similar* — *morally similar* and *physiologically similar* might not be as closely connected as this premise suggests (in any event, it needs further discussion).

One of the most famous instances of a dilemma and counter-dilemma is attributed to the sophist **Protagoras** (c.490-420 BCE) and his former student Euathlus. Protagoras had agreed to teach his student at a reduced fee, with the understanding that Euathlus would pay the remainder after winning his first court case. Some time passed, however, with Euathlus never taking on any cases, so Protagoras took him to court to sue for unpaid fees. Protagoras argued before the court — the Areopagus in Athens, where Socrates would later be tried — that Euathlus must pay him no matter what, offering the following dilemma: If Euathlus wins this court case, then (by the terms of his agreement with Protagoras), he will have to pay. And if he loses the case, then (by order of the court) he will have to pay. Since he will either win the court case or he won't, he will have to pay:

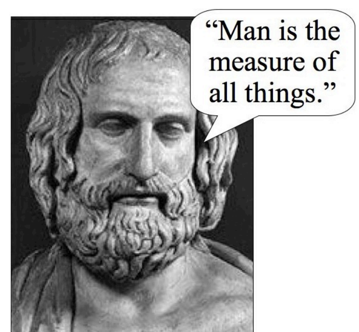

Protagoras
(Greece; c.490-c.420 BCE)

(1) **If p** (the student *loses* the court case), **then q** (he has to pay [by order of the court])
(2) **If not-p** (the student *wins* the court case), **then q** (he has to pay [by the terms of his agreement with Protagoras])
(3) **Either p or not p**
(4) Therefore, **q**

Euathlus had clearly learned his lessons well from Protagoras, however, and offered up a counter-dilemma: He will not have to pay Protagoras any fees, for if he wins the court case, then (by order of the court) he does not have to pay, and if he loses the court case, then (by the terms of his agreement with Protagoras) he will not have to pay. Since he will either win the court case or he won't, he will not have to pay:

(1) **If p** (the student wins the court case), **then q** (he doesn't have to pay [by order of the court])
(2) **If not-p** (the student loses the court case), **then q** (he doesn't have to pay [by the terms of his agreement with Protagoras])
(3) **Either p or not p**
(4) Therefore, **q**

We know of this court case from several ancient sources; these differ on various details, but they all agree that the court was unable to decide the case, given its paradoxical nature.

[7] INDUCTIVE LOGIC

All of the arguments discussed so far are *deductive*; the kind of inference used in those arguments guarantees the truth of the conclusion, so long as the premises are true. *Inductive* arguments, even good ones, do not provide conclusions that are certainly true; but good inductive arguments do provide conclusions that are more likely true than false, and therefore more reasonable to believe than not to believe. Because true premises do not guarantee a true conclusion in inductive arguments, we speak of good inductive arguments as being *cogent* (rather than sound), and the argument's form as being either *strong* or *weak* (rather than valid or invalid). A **strong argument** is one such that it is unlikely that its conclusion is false so long as its premises are true. A **cogent argument** is an inductive argument that is strong, and all of whose premises are true. There are many different kinds of inductive arguments; four of the more common are generalization from the past, argument from authority, argument from analogy, and hypothetical induction.

Generalization from the Past

Perhaps the most common form of inductive reasoning involves generalizing from past experiences; this is also called *enumerative induction*. In abstraction, it looks like this:

(1) Token 1 of type A has property X.
(2) Token 2 of type A has property X. [… and so on …]
(3) ∴ All tokens of type A have property X.

(where a *token* is an individual, and a *type* is a kind of individual, for instance, Socrates is a token of the type human). Here we draw a conclusion about some **target population** (all the tokens of type A) based on our observation of some **sample population** of that group (tokens 1, 2, ... of type A).

Now let's consider this kind of induction more concretely.

Fred is a beginning philosophy student cramming for his first exam of the semester. In preparation for an all-night study marathon, he heads to the local grocery store for supplies — carrot sticks, apples, and a large 24-count box of Twinkie snack cakes. "With such a combination of power food," Fred reasons, "I can't help but do well on tomorrow's exam." So now Fred is back in his room, trying to remember the difference between deductive and inductive logic while biting into his first Twinkie of the night — but something tastes odd and the Twinkie is crunchier than usual. He spits it out onto the pages of his opened philosophy text, and you can imagine his considerable surprise and disgust when he sees little brown objects mixed in with the white marshmallow filling. His roommate wanders by, considers the half-chewed mess on Fred's textbook, and offers helpfully: "Looks like you've got some mouse droppings in that Twinkie, Fred" — which is confirmed by a closer inspection with a magnifying glass. Fred can't bring himself to finish the half-eaten Twinkie, so he throws it away and opens another, but this second Twinkie tastes no better than the first, and it too ends up half-chewed on his textbook. Things continue in this manner until finally Fred reasons:

(1) The first Twinkie I ate from this box of 24 Twinkies had its crème filling riddled with mouse droppings.
(2) The second Twinkie I ate from the box was similarly contaminated.
(3) Likewise with the third Twinkie.
(4) ∴ All the Twinkies in this box have mouse droppings in their crème filling.

Now it's perfectly possible that the remainder of the Twinkies in the box are entirely free of mouse droppings — there's no guarantee that the next Twinkie will be like the previous Twinkie, nor that, more generally, the future will be like the past — but there is a strong presumption for it, and most people would reasonably throw out the remaining Twinkies (or better still, return them to the store for a refund). So the above argument is cogent.

We often engage in this kind of enumerative induction, but sometimes we do it badly. Faulty induction of this sort (commonly referred to as the *fallacy of hasty generalization*) results from either of two errors: (1) the sample size is too small for the group, or (2) the sample is not representative of the group. Consider the following two accounts of inductive reasoning:

(1) John held his hand directly over the flame of an acetylene torch for 5 seconds, and was badly burned. He concluded that such a flame would always burn human skin when placed in such close proximity.

(2) John wished to know what the professors at his university thought about Jack Kerouac as a novelist, so he asked his economics professor, who exclaimed that Kerouac was the greatest novelist of all time. From this John concluded that the professors at his university think that Jack Kerouac was the greatest novelist of all time.

[Extra]

Calculating the Margin of Error

Suppose I want to know how many white marbles are in a large box of five-million marbles, but I don't feel like counting all five-million marbles. Instead, I randomly draw 400 marbles from the box, and discover that 25% of this sample are white. How likely will it be that the entire population of five-million will have this proportion of white marbles?

The **margin of error** is a measure of this accuracy, and can be mathematically calculated and presented in tables like the one to the right (which is based on a 95% **level of significance**). The margin of error is found by dividing 1 by the square root of the number of individuals in the sample. So, for instance, the margin of error for a sample of 400 is exactly +/- 5%. The 95% level of significance means that if I were to take, say, 100 samples just like the one I took (returning the samples to the box each time), then at least 95 of the samples will contain white marbles in proportions ranging from 20% to 30%. This 10% gap is the **confidence interval**. If my sample size was only 100, then the confidence interval would be 20%.

Margin of Error	
Sample Size	Margin of Error
1600	2.5%
1000	3.2%
900	3.3%
800	3.5%
700	3.8%
600	4.1%
500	4.5%
400	5.0%
300	5.7%
200	7.1%
100	10.0%
50	14.1%

In both of these cases, John is reasoning from a sample of one to a much larger target population. In the first, the target population is quite large (namely, the skin of any human being anywhere); in the second, the target population consists of the 100 or so professors at his university. It would seem that his reasoning would be more reliable in the second example — since the sample size constitutes a larger proportion of the target population — but in fact the second example is less reliable, because the sample is much less representative of the group. With the first example, the target population is **homogenous** regarding the property being generalized (viz., sensitivity to a flame), and so a small sample size is adequate;[7] with the second example, however, there is likely considerable diversity regarding the property, and consequently a larger sample size would be needed to get a sense of how that property is related to the target population.

In this second example, because the property (viz., the degree to which Kerouac is admired as a novelist) is not homogenous across the general population, we must increase the sample size, so that

---

[7] This is a vexing issue. Inductive generalization is normally applied to groups of individuals where some property is arbitrarily distributed among them; we can then save time by looking at a sample of the group to arrive at a good estimate of how the property is related to that entire group. If this property is a necessary condition of being a member of that group, however — like having skin that is easily burned by an acetylene torch — then looking at a single individual is quite adequate for determining how the property applies to all the other individuals. Deciding whether a property is truly arbitrary in this sense, or is instead somehow connected with the nature of the individuals, is itself a complicated question. Do we discover this inductively?

In general, how do we know if a sample is representative of the target group? If I taste a single Twinkie from a box of 24, do I have any reason *not* to believe that all the Twinkies in that box will taste roughly the same? So having tasted one, can I appropriately infer that all the other Twinkies will taste the same? And if I taste mouse droppings in one, should I make this same inference?

our sample will more accurately reflect the general population regarding this property. For the same reason, we need to make sure that the sample is **random** (i.e., find some method by which it is just as likely that he choose one of them as any other). For instance, if we decided to increase the sample to 25 professors, but we drew them all from the campus Kerouac fan club, then the sample will still probably not be representative of the whole population.

Similarly, if we used a procedure that ensured randomness, but where the sample size was still too small, the sample would likely not reflect the larger population. Suppose there are the 100 professors at the university. John needs to devise some method for arriving at a **random sample**. Let John take 100 slips of papers, put the name of a different professor on each of these, place them all in a bag, give it a good shake, close his eyes, and then draw out a slip of paper. That should make the procedure random. John now reads the name of his economics professor, and discovers that this professor "greatly admires" Kerouac. From this random sampling, in which 100% of the sample population "greatly admires" Kerouac, John concludes that 100% of the target population also "greatly admires" Kerouac. This would be a mistaken inference, of course, because the sample size is too small to accurately represent the group. The sample was arrived at randomly, but it is too small to provide a workable margin of error (see below).

ARGUMENT FROM AUTHORITY

The vast majority of what we believe is based on hearsay. Someone whom we trust tells us that S is true, and so now we believe S. Very few of us have good independent reasons for believing in the existence of atoms, viruses, Pluto, the year of Napoleon's defeat at Waterloo, or even the true identity of our biological parents — save for the various authorities in whom we've placed our trust. There's clearly nothing wrong in any of this, either, so long as we're aware of what we're doing, and we take care in choosing our authorities.

All arguments from authority have the following structure:

> (1) A (some person) is an authority regarding S (some statement).
> (2) A believes S.
> (3) ∴ S is true.

In general, we would do well to listen to the experts. Plato, for instance, depicts Socrates as committed to the testimony of experts, those with genuine knowledge (*Crito*, 47a-e):

> *Socrates*: Consider then, do you not think it a good statement that one must not value all the opinions of men, but some and not others, nor the opinions of all men, but those of some and not of others? What do you say? Is this not well said?

Bertrand Russell, the 20th century British logician and philosopher, suggested that following just three basic rules would vastly improve our epistemic situation:

> (1) that when the experts are agreed, the opposite opinion cannot be held to be certain; (2) that when they are not agreed, no opinion can be regarded as certain by a non-expert; and (3) that when they all hold that no sufficient grounds for a positive opinion exist, the ordinary man would do well to suspend his judgment.[8]

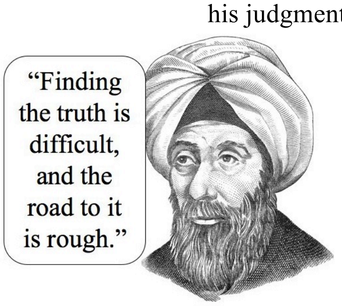

Ibn al-Haytham
(Basra/Cairo; 965-1039 C.E)

The basic problem facing this kind of argument is determining who is an authority. A common mistake is to accept as "an authority regarding P" someone whose authority lies in some area other than that of P. When individuals speak outside their area of expertise, their words should carry no more weight than the next person's. Authorities need not have higher degrees or any formal education at all; they simply need to be reliable sources of information of the relevant kind — they need to be honest and in a position to know. If the statement in question concerns the winner of the 1982 World Series, for instance, then any avid baseball fan will likely serve as a reliable authority.

---

[8] Bertrand Russell, *Let the People Think* (London: William Clowes, 1941), p. 2.

The matter gets more complicated as the statements become more contested or controversial, and individuals might harbor motives to lie or otherwise distort the truth. Here we typically find ourselves involved in a new inductive argument based on the person's past performance: Have they been reliable in the past? Have we ever discovered their claims to be false? This additional inductive step involves a "generalization from the past," for we are assuming that a person's past performance (as a reliable source of information) is indicative of their future performance. Finally, we must be sensitive to any special motivations that an otherwise reliable source might have to lie in the present case — for instance, if a lie would spare the person or those he cares about much harm or make possible some significant good.

ARGUMENT FROM ANALOGY

The first thing to note about arguments from analogy is how badly they can go. The second thing to note is how often we use them. First the bad, then the ubiquitous:

(1) Soccer and hockey share many similarities (e.g., they are both team sports that involve getting an object into a net).
(2) Hockey is played on ice.
(3) ∴ Soccer is probably played on ice as well.

Almost everything has something in common with everything else, so it would seem to be a rather sorry kind of reasoning to infer that two things, because they share *one* property, will also share some second property that we know the first thing to have. Nonetheless, and quite remarkably, we routinely make good use of this kind of argument.

Arguments from analogy all have the following form:

(1) Items A and B have characteristic X.
(2) Item A also has characteristic Y.
(3) ∴ B probably has characteristic Y as well.

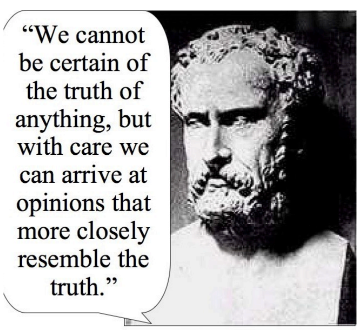

Xenophanes
(Colophon; 6th century B.C.E)

The thing or things we know best we call the **primary analogate(s)**. The thing about which we are, by way of analogy, drawing an inference, we call the **secondary analogate**. There might be only one primary analogate or there might be dozens; there might be only two characteristics being considered between the primary and secondary analogates, or there might be twenty or more. But all arguments from analogy compare at least two things sharing at least one characteristic, and from which we infer an additional characteristic in the secondary analogate that we *know* is possessed by the primary analogate.

In the U.S. legal system, judges routinely use analogy when deciding court cases based on precedent. Consider the case of *Adams v. New Jersey Steamboat Co.* (151 NY 163 [1896]): On June 17th, 1889, the plaintiff (Adams) boarded an overnight steamboat from New York City to Albany, and during the night money was stolen from his room. Adams sued the steamboat company for damages, noting that under the normal contract of hospitality, hotel proprietors are responsible for the safety of the guest's belongings. Because the steamboat was essentially operating like a "floating hotel," it should be bound by this well-established contract. The steamboat company, for its part, appealed to the decision given in *Carpenter v. Railroad Co.* (124 NY 53, 26 N. E. 277), in which the court found the railroad company *not* liable for luggage stolen from a sleeper berth, since the contract between the plaintiff and the defendant was primarily for travel and not for lodging. So, how should the judge decide the *Adams* case? Is the steamboat more like a "floating hotel" or like a "seagoing train"? Which similarities are more relevant? What are the disanalogies?

Consider the following analogical argument:

(1) Bill, Ted, and Al are all Manchester alumni.
(2) Bill and Ted both drive BMWs.
(3) ∴ Al probably also drives a BMW.

This is not a particularly strong argument as it stands, but there might be additional information that would strengthen the analogy. For instance, if Bill and Ted were picked at random from all Manchester alumni, then both of them driving BMWs might indicate a characteristic true of all or most of that population. If we increased the number of samples (by looking at the cars driven by other alumni), and found that all or most alumni sampled did in fact drive BMWs, then that would increase the likelihood that Al also drives a BMW. We also might increase the number of characteristics observed: suppose that Bill, Ted, and Al all majored in philosophy, graduated the same year, took similar jobs, etc. The more similarities we can find between the samples, the more likely the analogy will hold. But these similarities have to be *relevant*: noting that Bill, Ted, and Al all have first names with fewer than five letters does *not* make the analogy any stronger. And if we find that Bill and Ted were both pre-med majors and went on to become neurosurgeons, but Al majored in English and now writes poetry for a living, or that Bill and Ted graduated in 1955 and Al graduated just last year, then these are *disanalogies* which can weaken the analogy considerably (depending upon their relevance). In all, there are six different factors that determine the strength or weakness of an argument from analogy:

> **FINDING NEPTUNE**
>
> The planet Neptune was discovered in 1846. Here's how it happened. In 1845, a year before the discovery, the mathematician John Couch Adams (1819-1892) noticed perturbations in the orbit of Uranus that the current theory of celestial mechanics (viz., Newton's laws) did not allow. This was the surprising phenomenon in need of an explanation. Through abductive reasoning, Couch arrived at an hypothesis (H): *There is some undiscovered planet beyond Uranus causing the perturbations*. If H is correct, then we should be able to observe a new planet (i.e., a bright spot in the night sky) at a certain place in the sky at a certain time. On September 23, 1846, Johann Galle observed a bright spot at the predicted time and place; later observations showed it following the predicted path, thus confirming H (and further confirming Newtonian mechanics).

(1) **Relevance of similarities**: Relevance of the known shared property (x) to the inferred shared property (y). [The more relevant, the stronger the analogy.]

(2) **Disanalogy**: These are relevant differences between the primary and secondary analogates. [Usually, the more disanalogies, the weaker the analogy.]

(3) **Number of Similarities**: Number of similarities between primary and secondary analogates. [The more similarities, the stronger the analogy.]

(4) **Sample size**: Number and kind of primary analogates. [The more samples, the stronger the analogy.]

(5) **Sample diversity**: Diversity among the primary analogates; randomized sampling strengthens the likelihood of the secondary analogate sharing the contested property. [Usually, the greater the diversity, the stronger the analogy.]

(6) **Specificity**: Specificity of the conclusion relative to the premises. [The more specific the conclusion, the weaker the analogy.]

An argument from analogy with somewhat more historical importance was offered by **William Paley** (1743-1805), who argued that the universe and pocket watches share one important characteristic: They have many parts that seem to interact with one another towards some purpose. But since watches have the additional property of having been created by some "designing intelligence," we might conclude by analogy that the universe shares this property as well, and so was created by a designing intelligence, namely, a God as envisioned in the Judeo-Christian-Islamic tradition. The strength of this particular analogy will be discussed later.

As for *Adams v. New Jersey Steamboat Co.*, Judge O'Brien ruled in favor of Adams, writing that ...

> The relations that exist between a steamboat company and its passengers, who have procured staterooms for their comfort during the journey, differ in no essential respect from those that exist between the innkeeper and his guests. The passenger procures and pays for his room for the same reasons that a guest at an inn does. There are the same opportunities for fraud and plunder on the part of the carrier that was originally supposed to furnish a temptation to the landlord to violate his duty to the guest. A steamer carrying passengers upon the water, and furnishing them with rooms and entertainment is, for

> all practical purposes, a floating inn, and hence the duties which the proprietors owe to their charge ought to be the same. [...] The two relations, if not identical, bear such close analogy to each other that the same rule of responsibility should govern.[9]

HYPOTHETICAL INDUCTION

All arguments from hypothetical induction look something like the following:

(1) P [some phenomenon]
(2) If T, then P [T explains the existence/ occurrence of P]
(3) ... [Additional premises that speak against rival explanations of P]
(4) ∴ T is the best explanation we have of P.
(5) ∴ T

Formally, this kind of reasoning resembles the deductively fallacious *affirming the consequence*.

The father of American pragmatism, **Charles Sanders Peirce** (1839-1914), was the first to write explicitly about hypothetical induction, or what is often called "inference to the best explanation." Peirce called it "abductive logic," contrasting it with both deductive logic and enumerative forms of inductive logic:

> The great difference between induction [generalization] and hypothesis [abduction] is, that the former infers the existence of phenomena such as we have observed in cases which are similar, while hypothesis supposes something of a different kind from what we have directly observed, and frequently something which it would be impossible for us to observe directly.[10]

Peirce noted that we are remarkably successful with abduction; given the infinite number of possible hypotheses from which to choose, it is amazing how quickly we normally hit upon the right one:

> Think of what trillions of trillions of hypotheses might be made of which one only is true; and yet after two or three or at the very most a dozen guesses, the physicist hits pretty nearly on the correct hypothesis. By chance he would not have been likely to do so in the whole time that has elapsed since the earth was solidified.

Peirce understood the **scientific method** as a blending of abductive, deductive, and inductive reasoning. First, we begin with (1) a **surprising phenomenon**, which is in need of an explanation. This initiates the inquiry. We then use (2) **abductive reasoning** to develop an hypothesis [H] that explains the phenomenon. Next we use (3) **deductive reasoning** to determine one or more necessary conditions of H (that is, observable phenomena that will obtain under a certain condition [C], if H is true). Finally, we use (4) **inductive reasoning** to decide if the new observations offer adequate support for H (here, we are looking for the deduced necessary conditions of H, given C).[11] This is such an important, and common, kind of reasoning, it is worth looking at these steps in more detail.

The Case of Childbed Fever

Childbed fever (puerperal sepsis) has been a concern for women in childbirth since ancient times. Immediately after giving birth, the uterine wall where the placenta was attached is a large open wound and highly susceptible to infection. With childbed fever, the septic woman develops a high fever shortly after giving birth, and dies (at least without the intervention of antibiotic treatment — a 20th century development). As unfortunate as this was, it

---

[9] Quoted in Martin P. Golding, *Legal Reasoning* (Broadview Press, 2001), pp. 46-47.

[10] This and the next quoted passage come from Peirce's "Three Types of Reasoning" (1903). With enumerative induction, we infer from one individual to another, but it's always the same kind of individual. With abduction, on the other hand, the inference is normally to a completely different kind of thing, since it is what explains the observed individual. The inferred entity might even be impossible to observe directly (such as an electron, or the earth's core, or a black hole) — but it's existence and nature help explain phenomena that *are* directly observable.

[11] In "Three Types of Reasoning" (1903), Peirce wrote: "Induction consists in starting from a theory, deducing from it predictions of phenomena, and observing those phenomena in order to see *how nearly* they agree with the theory."

remained relatively rare and isolated until the 17th century advent of "laying-in" hospitals, where women would go to deliver their babies. The first recorded epidemic of childbed fever occurred in 1646 at the *Hôtel Dieu* in Paris, and reports of sporadic epidemics continued over the next two centuries on both sides of the Atlantic; an understanding of the disease, however, was slow in coming. Thomas Watson, a professor of medicine at King's College Hospital, London, suggested in 1842 that the fever was promoted by the attending physician having unwashed hands — in other words, that the fever was a matter of contagion, passed from woman to woman on the physician's hands — but his idea was ignored. The following year in Boston, Dr. Oliver Wendell Holmes also argued that childbed fever was caused by contagion. Apart from urging physicians to wash their hands with chlorinated water when moving from one woman to the next, Holmes also suggested that these physicians forego performing autopsies at the same time that they are attending women on the birthing ward. Holmes' ideas were met with ridicule; most physicians couldn't imagine that they were the actual cause of the problem.[12]

Unaware of the writings of Watson and Holmes, the Hungarian physician Ignaz Semmelweis was puzzling over a disturbing fact at the Vienna General Hospital where he worked. There were two delivery wards at the hospital — on the first ward, wealthier women were attended by physicians and their medical students, while on the second ward, poorer women were attended by midwives. The incidence of childbed fever was about 16% among the women in the first ward, but only 2% on the second ward. In hindsight, knowing what we know, this fact is not surprising, for the physicians and medical students were performing autopsies downstairs in the morgue while waiting for the babies to be born, and they weren't washing their hands carefully — so the same hand that had been probing the cavities of a dead body would the next minute be checking for cervical dilation.

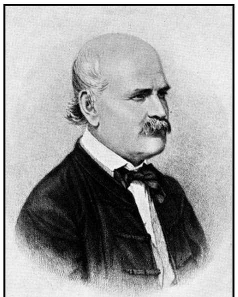
Ignaz Semmelweis
(Vienna, 1818-1865)

A sad epiphany came to Semmelweis when his friend Jakob Kolletschka, the professor of forensic pathology, accidentally cut himself while performing an autopsy, grew septic, and died; his symptoms were remarkably like that of a woman suffering from childbed fever. Semmelweis wrote:

> Suddenly a thought crossed my mind: childbed fever and the death of Professor Kolletschka were one and the same. His sepsis and childbed fever must originate from the same source ... the fingers and hands of students and doctors, soiled by recent dissections, carry those death dealing cadavers' poisons into the genital organs of women in childbirth ....

Working with this hunch, Semmelweis had all his physicians and students wash their hands in chlorinated water before attending the laboring women, and with this simple practice the incidence of childbed fever dropped to less than 3%. Despite these dramatic results, his ideas were dismissed, and many more women died unnecessarily before the medical community accepted the idea that childbed fever was the direct result of physician-caused contagion.

We can reconstruct Semmelweis' reasoning by patterning it after Peirce's model of the scientific method: (1) a **surprising phenomenon** [P] in need of an explanation, (2) **abductive reasoning** to develop an hypothesis [H] that explains the phenomenon, (3) **deductive reasoning** to determine one or more necessary conditions [NC] of H (if H is true, then necessary conditions of H will also have to be true), and (4) **inductive reasoning** to test for these necessary conditions.

(1) Phenomenon (P): We observe some phenomenon that is surprising or otherwise inexplicable, given the current laws and theories. This is the *explanandum*, in need of an *explanans*. In the above case:

P = Five times more women contract childbed fever if they are on the ward attended by the physicians rather than by midwives.

(2) Hypothesis (H): There are always many ways of explaining why P might happen; the goal is to find the explanation that is most likely correct. Just like P, this explanatory hypothesis (H) should be stated as clearly as possible. Semmelweis' hypothesis was something like this:

---

[12] If you find such obtuseness among physicians inconceivable, remember that the germ theory of disease (viz., the theory that many diseases are caused by microscopically small organisms — what we today call bacteria and viruses) wasn't well established until 1875, with the work of Robert Koch in Germany.

H = The physicians are transmitting the cause of this fever (what Semmelweis called "cadaveric matter") on their hands, which they pick up during their autopsies in the morgue.

(3) Necessary Condition (NC) of H: Explanations can be thought of as the antecedent clause of a conditional statement, where the thing to be explained is the consequent clause. So: **If H, then P**. H is the sufficient condition of P, and P the necessary condition of H.[13] We know that P is true (it's the surprising phenomenon that we want to explain), but we don't know if H is true. So we need to test if H is true, and one way to do this is to find some *second* necessary condition of H, that is, some new conditional statement: **if H, then NC**. Once we do that, we can check to see if NC is true; if NC is false, then by *modus tollens* we know that NC is also false. If NC is true, then NC has not been falsified, and so might still be true. This process "confirms" the truth of H, but never can prove it.[14]

This step is more abstract than the others, so let me repeat the above, with examples.

If the hypothesis (H) is to count as an explanation of P, then P has to be a necessary condition of H (if H, then P). An explanation is always a sufficient condition of the thing being explained. Our problem is that, for any particular phenomenon, there might be many possible explanations, and we want to discover the correct one. For instance, you hear a dog yelp in pain. One explanation of the yelping is that someone kicked the dog, but there are other possible explanations: it was stung by a bee, it was chewing on an electrical cord and received a shock, it has gas pains, and so on. Given the many possible explanations, we need to (abductively) arrive at what we think is the most plausible explanation, and then we need to (deductively) devise a test for it. In other words, we need to arrive at some *other* necessary condition of H, for which we can then (inductively) test. If it passes this test, then the presence or truth of H is confirmed — although not proven — making it available as a possible cause or proper explanation of P. For Semmelweis, the additional necessary condition was:

NC = If transmission of this cadaveric matter to the pregnant women is stopped, then the additional illnesses should disappear.

(4) Testing for the NC: We now need to devise a test to see if the new NC is present.

Test = Have all attending physicians and medical students wash their hands with chlorinated water, which should destroy the cadaveric matter.

Observed phenomenon (OP) under the test condition: This is the observation of the presence or absence (or the truth/falsity) of NC.

OP = Incidence of childbed fever drops dramatically.

Conclusion: The explanatory hypothesis (H) is confirmed or falsified. If confirmed, we generally refer to the hypothesis as a theory.

Criteria for Deciding between Theories

Hypotheses are supported, but never proven as true, by the inductive reasoning in stage four. Galle's astronomical observations confirmed the existence of a new planet and served as a remarkable confirmation of Newtonian mechanics, but it certainly didn't prove that Newton was right — and we have since learned that Newton was, in fact, wrong (his mechanics have been replaced by Einstein's theories of relativity as a more adequate account of the universe). Theories are kept until they are replaced by something more adequate. The criteria for adequacy, or for preferring one theory over another, are something like the following:[15]

(1) **Internally consistent**: The theory involves no self-contradiction.

(2) **Testable/Predictive**: The theory must be something that we can test by observation; more specifically, it must be falsifiable (we must be able to indicate the conditions that would falsify the theory: for this to be possible, the theory must make some new prediction that can be checked).

(3) **Externally consistent**: The theory is consistent with well-established facts and other theories.

---

[13] For any conditional statement — if p, then q — p is a sufficient condition of q, and q is a necessary condition of p. If this isn't clear, see the discussion of necessary and sufficient conditions, above.

[14] If we could prove H in this deductive fashion, then natural science would be a matter of deductive reasoning, and would enjoy the same kind of certainty as mathematics; but natural science is about the physical world — see the quoted passage from Stephen Jay Gould, below — and so will always remain open-ended and uncertain.

[15] See Thomas Kuhn, *The Essential Tension* (1977), pp. 321-22.

(4) **Simple**: The theory introduces no unnecessary mysteries (Ockham's Razor).

(5) **Scope**: The more that the hypothesis can explain, the more it is to be preferred. Newton's account of motion had greater scope than Galileo's. Einstein's account had greater scope than Newton's.

(6) **Fruitful**: The theory suggests various other lines of inquiry.

Facts, Laws, and Theories

This is a good place to briefly consider a common confusion among non-scientists regarding the relationship between facts and theories. The paleontologist Stephen Jay Gould (1941-2002) offers a helpful clarification in the context of his own field, evolutionary biology:

> "By doubting, we come to inquiry; and through inquiry, we perceive truth."
> — Peter Abelard (1079-1142)

> In the American vernacular, "theory" often means "imperfect fact" — part of a hierarchy of confidence running downhill from *fact* to *theory* to *hypothesis* to *guess*. Thus the power of the creationist argument: evolu-tion is "only" a theory and intense debate now rages about many aspects of the theory. If evolution is worse than a fact, and scientists can't even make up their minds about the theory, then what confidence can we have in it? Indeed, President Reagan echoed this argument before an evangelical group in Dallas when he said (in what I devoutly hope was campaign rhetoric): "Well, it is a theory. It is a scientific theory only, and it has in recent years been challenged in the world of science — that is, not believed in the scientific community to be as infallible as it once was."
>
> Well evolution *is* a theory. It is also a fact. And facts and theories are different things, not rungs in a hierarchy of increasing certainty. Facts are the world's data. Theories are structures of ideas that explain and interpret facts. Facts don't go away when scientists debate rival theories to explain them. Einstein's theory of gravitation replaced Newton's in this century, but apples didn't suspend themselves in midair, pending the outcome. And humans evolved from ape-like ancestors whether they did so by Darwin's proposed mechanism or by some other yet to be discovered.
>
> Moreover, "fact" doesn't mean "absolute certainty"; there ain't no such animal in an exciting and complex world. The final proofs of logic and mathematics flow deductively from stated premises and achieve certainty only because they are *not* about the empirical world. Evolutionists make no claim for perpetual truth, though creationists often do (and then attack us falsely for a style of argument that they themselves favor). In science "fact" can only mean "confirmed to such a degree that it would be perverse to withhold provisional consent." I suppose that apples might start to rise tomorrow, but the possibility does not merit equal time in physics classrooms.
>
> Evolutionists have been very clear about this distinction of fact and theory from the very beginning, if only because we have always acknowledged how far we are from completely understanding the mechanisms (theory) by which evolution (fact) occurred. Darwin continually emphasized the difference between his two great and separate accomplishments: establishing the fact of evolution, and proposing a theory — natural selection — to explain the mechanism of evolution.[16]

Facts are observations, while **theories** are explanations of what is observed. An especially useful kind of fact is a **law**, which is a description of a regularity found among the phenomena. The claim that "all swans are white" is a law (and, as we now know, a false law), for it describes how certain phenomena (being a swan, being white) combine in our experience. Boyles' Law (formulated in 1662 by the Irish chemist Robert Boyle) describes how the pressure and volume of a gas are related (viz., they vary inversely), while Charles' Law describes the relationship between the volume and temperature of gases, and Gay-Lussac's Law describes the relationship between temperature and pressure. These laws all describe the way the gases behave, and as such are quite different from the *kinetic theory of gases* — namely, that a volume of gas consists of individual gas molecules colliding with each other and with the sides of the container holding the gas, and that it is the collection of these collisions that accounts for the observed temperature, pressure, and volume of the gas. This kinetic theory (first postulated by Daniel Bernoulli in 1740, and in opposition to the Newtonians) explains why these law-like regularities obtain. Similarly, Newton's inverse-square law of gravity is an excellent example of a law: it describes a regularity in the observed motions of bodies, but leaves wholly unexplained the mechanism behind this regularity (i.e., what gravity is).

---

[16] Stephen J. Gould, "Evolution as Fact and Theory," *Discover* (May 1981).

Another example: We can make various astronomical observations of the position of the sun, moon, planets, and stars. Each observation of the sun, for example, would count as a **fact**: "The sun is at position, p, and time, t." We accumulate these many facts and notice clear regularities, which can eventually be distilled into **laws**, such as: "The sun returns to the same place in the sky every 23 hours and 56 minutes" or "The observed location of the sun follows the same path — which we call the ecliptic — each day." These laws are nothing more than summaries of the many observations (although they go beyond the observations, of course; they predict what we will observe in the future, and since our set of past observations is always incomplete, it claims what we would have observed, had we looked at that time and place.) A **theory** will explain why these facts and laws are as they are, and here we have two well known and competing theories. The first theory (geocentrism) is that the earth is motionless while the sun moves around the earth approximately every 24 hours; that the sun literally rises in the morning, moves westward across the sky, and then sets in the evening. The second theory (heliocentrism) is that the earth spins on its north-south axis once approximately every 24 hours, and that it is this spinning that causes the apparent motion of the sun each day. These two theories offer quite different explanations for our many astronomical observations (facts).

Laws often act like theories, because we often use them as a kind of explanation: "He's flatulent because of all that chili he ate for dinner" — this is an explanation based on the law-like correlation between eating chili and suffering flatulence. There is a deeper explanation, however, based on the chemistry of beans and digestive tracts, and here we are at the level of theoretical explanation. The difference between these two levels of explanation, and between laws and theories, is complicated and contested, but one useful way of thinking about it is this: With a law, one explains that a particular S is P by pointing to the law that "All S are P" (or that, e.g., 89% of all S's are P) ; with a theory, one explains *why* all S are P (based on the presumed natures of S and P).

INDUCTION, DEDUCTION, AND UNCERTAINTY

Inductive arguments have conclusions that are uncertain, but so do many deductive arguments. The difference is that, with deductive arguments, the uncertainty stems from one or more uncertain premises, while with inductive arguments the uncertainty is (at least in part) a result of the inference from the premises to the conclusion. For instance, all conclusions about the future are uncertain, because the future is itself uncertain. Inductive generalizations typically involve some claim about a future event, e.g., about the next swan that you examine (or more ambitiously, about all future swans). This conclusion was probably inferred from premises that were themselves wholly certain (e.g., I look at one swan and see that it is white, I look at a second swan and see that it is white, and so forth):

> This swan is white.
> That swan is white. [...]
> ∴ All swans are white.

But consider the following deductive argument:

> All swans are white.
> Alice has a swan.
> ∴ Alice's swan is white.

The conclusion is guaranteed to be true if the premises are both true, yet we could imagine the conclusion being false (since Alice might be in possession of a black swan). This uncertainty in the conclusion is not a result of the inference, but rather of an uncertainty found in the first premise.