

Critical Thinking

Critical Thinking

BRIAN KIM

OKLAHOMA STATE UNIVERSITY LIBRARIES
STILLWATER, OK



Critical Thinking by Brian Kim is licensed under a Creative Commons Attribution 4.0 International License, except where otherwise noted.

Contents

Attribution	vii
About this Resource	viii
1. Introduction to Critical Thinking	1
2. Logic and the Study of Arguments	11
3. Fallacies	34
4. Sentential Logic	49
5. Truth Tables	77
6. Categorical Logic	94

Attribution

Original content in this work is licensed CC-BY unless otherwise noted; links and non-original content maintain their original licensing. Footnotes are included throughout which point to the original works as they were available at the time this work was written.

Portions of this book, Chapters 4 and 5 in particular, have been incorporated from forallx: An Introduction to Formal Logic by P. D. Magnus, licensed CC-BY. Dr. Magnus' work was originally released in 2005, and he has shared the project github repository at this link.

About this Resource

Other file formats available

In addition to the web version, this book is available in a number of file formats including PDF, EPUB (for eReaders), MOBI (for Kindles), and various editable files. Here is a link to where you can download this book in another file format. Look for the Download this book drop-down menu to select the file type you want.

Critical Thinking by Brian Kim has been shared on the sites listed below. If you are using this book please email kathy.essmiller@okstate.edu. We would love to know so that we can keep you up to date on changes or collaborate with you on future iterations.

OERCommons

Let us know if you are having problems accessing this resource

We are always looking for ways to make our resources more accessible. If you have problems accessing this resource, please contact us to let us know, so we can fix the issue.

Please include the following information:

- The location of the problem by providing a web address or page description
- A description of the problem
- The computer, software, browser, and any assistive technology you are using that can help us diagnose and solve your issue

- e.g., Windows 10, Google Chrome (Version 65.0.3325.181), NVDA screen reader

You can contact us by completing this web form or by sending an email to kathy.essmiller@okstate.edu.

I. Introduction to Critical Thinking

I. What is Critical Thinking?^I

Critical thinking is the ability to think clearly and rationally about what to do or what to believe. It includes the ability to engage in reflective and independent thinking. Someone with critical thinking skills is able to do the following:

- Understand the logical connections between ideas.
- Identify, construct, and evaluate arguments.
- Detect inconsistencies and common mistakes in reasoning.
- Solve problems systematically.
- Identify the relevance and importance of ideas.
- Reflect on the justification of one's own beliefs and values.

Critical thinking is not simply a matter of accumulating information. A person with a good memory and who knows a lot of facts is not necessarily good at critical thinking. Critical thinkers are able to deduce consequences from what they know, make use of information to solve problems, and to seek relevant sources of information to inform themselves.

Critical thinking should not be confused with being argumentative or being critical of other people. Although critical thinking skills can be used in exposing fallacies and bad reasoning, critical thinking can also play an important role in cooperative reasoning and constructive tasks. Critical thinking can help us acquire knowledge, improve our theories, and strengthen arguments. We can also use critical thinking to enhance work processes and improve social institutions.

Some people believe that critical thinking hinders creativity because critical thinking requires following the rules of logic and rationality, whereas creativity might require breaking those rules. This is a misconception. Critical thinking is quite compatible with thinking “out-of-the-box,” challenging consensus views, and pursuing less popular approaches. If anything, critical thinking is an essential part of creativity because we need critical thinking to evaluate and improve our creative ideas.

II. The Importance of Critical Thinking

Critical thinking is a domain-general thinking skill. The ability to think clearly and rationally is important whatever we choose to do. If you work in education, research, finance, management or the legal profession, then critical thinking is obviously important. But critical thinking skills are not restricted to a particular subject area. Being able to think well and solve problems systematically is an asset for any career.

Critical thinking is very important in the new knowledge economy. The global knowledge economy is driven by information and technology. One has to be able to deal with changes quickly and effectively. The new economy places increasing demands on flexible intellectual skills, and the ability to analyze information and integrate diverse sources of knowledge in solving problems. Good critical thinking promotes such thinking skills, and is very important in the fast-changing workplace.

Critical thinking enhances language and presentation skills. Thinking clearly and systematically can improve the way we express our ideas. In learning how to analyze the logical structure of texts, critical thinking also improves comprehension abilities.

Critical thinking promotes creativity. To come up with a creative solution to a problem involves not just having new ideas. It must

also be the case that the new ideas being generated are useful and relevant to the task at hand. Critical thinking plays a crucial role in evaluating new ideas, selecting the best ones and modifying them if necessary.

Critical thinking is crucial for self-reflection. In order to live a meaningful life and to structure our lives accordingly, we need to justify and reflect on our values and decisions. Critical thinking provides the tools for this process of self-evaluation.

Good critical thinking is the foundation of science and democracy. Science requires the critical use of reason in experimentation and theory confirmation. The proper functioning of a liberal democracy requires citizens who can think critically about social issues to inform their judgments about proper governance and to overcome biases and prejudice.

Critical thinking is a metacognitive skill. What this means is that it is a higher-level cognitive skill that involves thinking about thinking. We have to be aware of the good principles of reasoning, and be reflective about our own reasoning. In addition, we often need to make a conscious effort to improve ourselves, avoid biases, and maintain objectivity. This is notoriously hard to do. We are all able to think but to think well often requires a long period of training. The mastery of critical thinking is similar to the mastery of many other skills. There are three important components: theory, practice, and attitude.

III. Improving Our Thinking Skills

Theory

If we want to think correctly, we need to follow the correct rules of reasoning. Knowledge of theory includes knowledge of these rules.

These are the basic principles of critical thinking, such as the laws of logic, and the methods of scientific reasoning, etc.

Also, it would be useful to know something about what not to do if we want to reason correctly. This means we should have some basic knowledge of the mistakes that people make. First, this requires some knowledge of typical fallacies. Second, psychologists have discovered persistent biases and limitations in human reasoning. An awareness of these empirical findings will alert us to potential problems.

Practice

However, merely knowing the principles that distinguish good and bad reasoning is not enough. We might study in the classroom about how to swim, and learn about the basic theory, such as the fact that one should not breathe underwater. But unless we can apply such theoretical knowledge through constant practice, we might not actually be able to swim.

Similarly, to be good at critical thinking skills it is necessary to internalize the theoretical principles so that we can actually apply them in daily life. There are at least two ways to do this. One is to perform lots of quality exercises. These exercises don't just include practicing in the classroom or receiving tutorials; they also include engaging in discussions and debates with other people in our daily lives, where the principles of critical thinking can be applied. The second method is to think more deeply about the principles that we have acquired. In the human mind, memory and understanding are acquired through making connections between ideas.

Attitudes

Good critical thinking skills require more than just knowledge and practice. Persistent practice can bring about improvements only if one has the right kind of motivation and attitude. The following attitudes are not uncommon, but they are obstacles to critical thinking:

- I prefer being given the correct answers rather than figuring them out myself.
- I don't like to think a lot about my decisions as I rely only on gut feelings.
- I don't usually review the mistakes I have made.
- I don't like to be criticized.

To improve our thinking we have to recognize the importance of reflecting on the reasons for belief and action. We should also be willing to engage in debate, break old habits, and deal with linguistic complexities and abstract concepts.

The *California Critical Thinking Disposition Inventory* is a psychological test that is used to measure whether people are disposed to think critically. It measures the seven different thinking habits listed below, and it is useful to ask ourselves to what extent they describe the way we think:

1. Truth-Seeking—Do you try to understand how things really are? Are you interested in finding out the truth?
2. Open-Mindedness—How receptive are you to new ideas, even when you do not intuitively agree with them? Do you give new concepts a fair hearing?
3. Analyticity—Do you try to understand the reasons behind things? Do you act impulsively or do you evaluate the pros and cons of your decisions?
4. Systematicity—Are you systematic in your thinking? Do you break down a complex problem into parts?

5. Confidence in Reasoning—Do you always defer to other people? How confident are you in your own judgment? Do you have reasons for your confidence? Do you have a way to evaluate your own thinking?
6. Inquisitiveness—Are you curious about unfamiliar topics and resolving complicated problems? Will you chase down an answer until you find it?
7. Maturity of Judgment—Do you jump to conclusions? Do you try to see things from different perspectives? Do you take other people's experiences into account?

Finally, as mentioned earlier, psychologists have discovered over the years that human reasoning can be easily affected by a variety of cognitive biases. For example, people tend to be over-confident of their abilities and focus too much on evidence that supports their pre-existing opinions. We should be alert to these biases in our attitudes towards our own thinking.

IV. Defining Critical Thinking

There are many different definitions of critical thinking. Here we list some of the well-known ones. You might notice that they all emphasize the importance of clarity and rationality. Here we will look at some well-known definitions in chronological order.

1) Many people trace the importance of critical thinking in education to the early twentieth-century American philosopher John Dewey. But Dewey did not make very extensive use of the term “critical thinking.” Instead, in his book *How We Think* (1910), he argued for the importance of what he called “reflective thinking”:

...[when] the ground or basis for a belief is deliberately sought and its adequacy to support the belief examined. This process is called reflective thought; it alone is truly educative in value...

Active, persistent and careful consideration of any belief or supposed form of knowledge in light of the grounds that support it, and the further conclusions to which it tends, constitutes reflective thought.

There is however one passage from *How We Think* where Dewey explicitly uses the term “critical thinking”:

The essence of critical thinking is suspended judgment; and the essence of this suspense is inquiry to determine the nature of the problem before proceeding to attempts at its solution. This, more than any other thing, transforms mere inference into tested inference, suggested conclusions into proof.

2) The *Watson-Glaser Critical Thinking Appraisal* (1980) is a well-known psychological test of critical thinking ability. The authors of this test define critical thinking as:

...a composite of attitudes, knowledge and skills. This composite includes: (1) attitudes of inquiry that involve an ability to recognize the existence of problems and an acceptance of the general need for evidence in support of what is asserted to be true; (2) knowledge of the nature of valid inferences, abstractions, and generalizations in which the weight or accuracy of different kinds of evidence are logically determined; and (3) skills in employing and applying the above attitudes and knowledge.

3) A very well-known and influential definition of critical thinking comes from philosopher and professor Robert Ennis in his work “A Taxonomy of Critical Thinking Dispositions and Abilities” (1987):

Critical thinking is reasonable reflective thinking that is focused on deciding what to believe or do.

4) The following definition comes from a statement written in 1987 by the philosophers Michael Scriven and Richard Paul for the National Council for Excellence in Critical Thinking ([link](#)), an organization promoting critical thinking in the US:

Critical thinking is the intellectually disciplined process of actively and skillfully conceptualizing, applying,

analyzing, synthesizing, and/or evaluating information gathered from, or generated by, observation, experience, reflection, reasoning, or communication, as a guide to belief and action. In its exemplary form, it is based on universal intellectual values that transcend subject matter divisions: clarity, accuracy, precision, consistency, relevance, sound evidence, good reasons, depth, breadth, and fairness. It entails the examination of those structures or elements of thought implicit in all reasoning: purpose, problem, or question-at-issue, assumptions, concepts, empirical grounding; reasoning leading to conclusions, implications and consequences, objections from alternative viewpoints, and frame of reference.

The following excerpt from Peter A. Facione's "Critical Thinking: A Statement of Expert Consensus for Purposes of Educational Assessment and Instruction" (1990) is quoted from a report written for the American Philosophical Association:

We understand critical thinking to be purposeful, self-regulatory judgment which results in interpretation, analysis, evaluation, and inference, as well as explanation of the evidential, conceptual, methodological, criteriological, or contextual considerations upon which that judgment is based. CT is essential as a tool of inquiry. As such, CT is a liberating force in education and a powerful resource in one's personal and civic life. While not synonymous with good thinking, CT is a pervasive and self-rectifying human phenomenon. The ideal critical thinker is habitually inquisitive, well-informed, trustful of reason, open-minded, flexible, fairminded in evaluation, honest in facing personal biases, prudent in making judgments, willing to reconsider, clear about issues, orderly in complex matters, diligent in seeking relevant information, reasonable in the selection of criteria, focused in inquiry, and persistent in seeking results which are as precise as the subject and the circumstances of inquiry permit. Thus, educating good

critical thinkers means working toward this ideal. It combines developing CT skills with nurturing those dispositions which consistently yield useful insights and which are the basis of a rational and democratic society.

V. Two Features of Critical Thinking

A. How? not What?

Critical thinking is concerned not with *what* you believe, but rather *how* or *why* you believe it. Most classes, such as those on biology or chemistry, teach you what to believe about a subject matter. In contrast, critical thinking is not particularly interested in what the world is, in fact, like. Rather, critical thinking will teach you how to form beliefs and how to think. It is interested in the type of reasoning you use when you form your beliefs, and concerns itself with whether you have good reasons to believe what you believe. Therefore, this class isn't a class on the psychology of reasoning, which brings us to the second important feature of critical thinking.

B. Ought Not Is (or Normative Not Descriptive)

There is a difference between normative and descriptive theories. Descriptive theories, such as those provided by physics, provide a picture of how the world factually behaves and operates. In contrast, normative theories, such as those provided by ethics or

political philosophy, provide a picture of how the world *should* be. Rather than ask question such as why something is the way it is, normative theories ask how something should be. In this course, we will be interested in normative theories that govern our thinking and reasoning. Therefore, we will not be interested in how we actually reason, but rather focus on how we ought to reason.

In the introduction to this course we considered a selection task with cards that must be flipped in order to check the validity of a rule. We noted that many people fail to identify all the cards required to check the rule. This is how people do in fact reason (descriptive). We then noted that you must flip over two cards. This is how people ought to reason (normative).



Critical Thinking by Brian Kim is licensed under a Creative Commons Attribution 4.0 International License, except where otherwise noted.

Notes

1. Section I-IV are taken from <http://philosophy.hku.hk/think/> and are in use under the creative commons license. Some modifications have been made to the original content.

2. Logic and the Study of Arguments

If we want to study how *we ought to reason* (normative) we should start by looking at the primary way that *we do reason* (descriptive): through the use of arguments. In order to develop a theory of good reasoning, we will start with an account of what an argument is and then proceed to talk about what constitutes a “good” argument.

I. Arguments

- Arguments are a set of statements (premises and conclusion).
- The premises provide evidence, reasons, and grounds for the conclusion.
- The conclusion is what is being argued for.
- An argument attempts to draw some logical connection between the premises and the conclusion.
- And in doing so, the argument expresses an inference: a process of reasoning from the truth of the premises to the truth of the conclusion.

Premises: The premises (and there can be more than one) are the statements being offered in support for the conclusion. The premises also embody the reasons or facts providing evidence for the conclusion’s credibility.

Conclusion: The conclusion is the statement being argued for.

Example: The world will end on August 6, 2045. I know this because my dad told me so and my dad is smart.

In this instance, the conclusion is the first sentence (“The world will end...”); the premises (however dubious) are revealed in the second sentence (“I know this because...”).

II. Statements

Conclusions and premises are articulated in the form of *statements*. Statements are sentences that can be determined to possess or lack truth. Some examples of true-or-false statements can be found below. (Notice that while some statements are categorically true or false, others may or may not be true depending on *when* they are made or *who* is making them.)

Examples of sentences that are statements:

1. It is below 40°F outside.
2. Oklahoma is north of Texas.
3. The Denver Broncos will make it to the Super Bowl.
4. Russell Westbrook is the best point guard in the league.
5. I like broccoli.
6. I shouldn't eat French fries.
7. Time travel is possible.
8. If time travel is possible, then you can be your own father or mother.

However, there are many sentences that cannot so easily be determined to be true or false. For this reason, these sentences identified below are not considered statements.

1. Questions: “What time is it?”
2. Commands: “Do your homework.”
3. Requests: “Please clean the kitchen.”
4. Proposals: “Let’s go to the museum tomorrow.”

Question: Why are arguments only made up of statements?

First, we only believe statements. It doesn’t make sense to talk about believing questions, commands, requests or proposals. Contrast sentences on the left that are not statements with sentences on the right that are statements:

Non-statements	Statements
What time is it? Do your homework.	The time is 11:00 a.m. My teacher wants me to do my homework.

It would be non-sensical to say that we believe the non-statements (e.g. “I believe what time is it?”). But it makes perfect sense to say that we believe the statements (e.g. “I believe the time is 11 a.m.”). If conclusions are the statements being argued for, then they are also ideas we are being persuaded to believe. **Therefore, only statements can be conclusions.**

Second, only statements can provide reasons to believe.

- Q: Why should I believe that it is 11:00 a.m.? A: Because the clock says it is 11a.m.
- Q: Why should I believe that we are going to the museum tomorrow? A: Because today we are making plans to go.

Sentences that cannot be true or false cannot provide reasons to believe. So, if premises are meant to provide reasons to believe, then only statements can be premises.

III. Representing Arguments

As we concern ourselves with arguments, we will want to represent our arguments in some way, indicating which statements are the premises and which statement is the conclusion. We shall represent arguments in two ways. For both ways, we will number the premises.

In order to identify the conclusion, we will either label the conclusion with a (c) or (conclusion). Or we will mark the conclusion with the \therefore symbol

Example Argument:

There will be a war in the next year. I know this because there has been a massive buildup in weapons. And every time there is a massive buildup in weapons, there is a war. My guru said the world will end on August 6, 2045.

- There has been a massive buildup in weapons.
- Every time there has been a massive buildup in weapons, there is a war.

(c) There will be a war in the next year.

Or

- There has been a massive buildup in weapons.
- Every time there has been a massive buildup in weapons, there is a war.

\therefore There will be a war in the next year.

Of course, arguments do not come labeled as such. And so we must be able to look at a passage and identify whether the passage contains an argument and if it does, we should also be identify which statements are the premises and which statement is the conclusion. This is harder than you might think!

There was a massive stampede outside of Tulsa. Crops were trampled and some cows were killed. Officials are not sure why the stampede started, but they are in the process of investigating the cause.

Example:

There is no argument here. There is no statement being argued for. There are no statements being used as reasons to believe. This is simply a report of information.

The following are also not arguments:

Advice: Be good to your friends; your friends will be good to you.

Warnings: No lifeguard on duty. Be careful.

Associated claims: Fear leads to anger. Anger leads to the dark side.

When you have an argument, the passage will express some process of reasoning. There will be statements presented that serve to help the speaker building a case for the conclusion.

IV. How to Look for Arguments¹

How do we identify arguments in real life? There are no easy, mechanical rules, and we usually have to rely on the context in order to determine which are the premises and the conclusions. But sometimes the job can be made easier by the presence of certain

premise or conclusion indicators. For example, if a person makes a statement, and then adds “this is because ...,” then it is quite likely that the first statement is presented as a conclusion, supported by the statements that come afterward. Other words in English that might be used to indicate the premises to follow include:

- since
- firstly, secondly, ...
- for, as, after all
- assuming that, in view of the fact that
- follows from, as shown / indicated by
- may be inferred / deduced / derived from

Of course whether such words are used to indicate premises or not depends on the context. For example, “since” has a very different function in a statement like “I have been here since noon,” unlike “X is an even number since X is divisible by 4.” In the first instance (“since noon”) “since” means “from.” In the second instance, “since” means “because.”

Conclusions, on the other hand, are often preceded by words like:

- therefore, so, it follows that
- hence, consequently
- suggests / proves / demonstrates that
- entails, implies

Here are some examples of passages that do not contain arguments.

1. When people sweat a lot they tend to drink more water. [Just a single statement, not enough to make an argument.]

2. Once upon a time there was a prince and a princess. They lived happily together and one day they decided to have a baby. But the baby grew up to be a nasty and cruel person and they regret it very much. [A chronological description of facts composed of statements but no premise or conclusion.]

3. Can you come to the meeting tomorrow? [A question that does not contain an argument.]

Exercises

Do these passages contain arguments? If so, what are their conclusions?

1. Cutting the interest rate will have no effect on the stock market this time around, as people have been expecting a rate cut all along. This factor has already been reflected in the market.
2. So it is raining heavily and this building might collapse. But I don't really care.
3. Virgin would then dominate the rail system. Is that something the government should worry about? Not necessarily. The industry is regulated, and one powerful company might at least offer a more coherent schedule of services than the present arrangement has produced. The reason the industry was broken up into more than 100 companies at privatization was not operational, but political: the Conservative government thought it would thus be harder to renationalize (The Economist 12/16/2000).
4. Bill will pay the ransom. After all, he loves his wife and children and would do everything to save them.
5. All of Russia's problems of human rights and democracy come back to three things: the legislature, the executive and the judiciary. None works as well as it should. Parliament passes laws in a hurry, and has neither the ability nor the will to call high officials to account. State officials abuse human rights (either on their own, or on orders from on high) and work with remarkable slowness and disorganization. The courts almost completely fail in their role as the ultimate safeguard of freedom and order (The Economist 11/25/2000).
6. Most mornings, Park Chang Woo arrives at a train station in central Seoul, South Korea's capital. But he is not commuter.

He is unemployed and goes there to kill time. Around him, dozens of jobless people pass their days drinking soju, a local version of vodka. For the moment, middle-aged Mr. Park would rather read a newspaper. He used to be a bricklayer for a small construction company in Pusan, a southern port city. But three years ago the country's financial crisis cost him that job, so he came to Seoul, leaving his wife and two children behind. Still looking for work, he has little hope of going home any time soon (The Economist 11/25/2000).

7. For a long time, astronomers suspected that Europa, one of Jupiter's many moons, might harbour a watery ocean beneath its ice-covered surface. They were right. Now the technique used earlier this year to demonstrate the existence of the European ocean has been employed to detect an ocean on another Jovian satellite, Ganymede, according to work announced at the recent American Geo-physical Union meeting in San Francisco (The Economist 12/16/2000).
8. There are no hard numbers, but the evidence from Asia's expatriate community is unequivocal. Three years after its handover from Britain to China, Hong Kong is unlearning English. The city's gweilos (Cantonese for "ghost men") must go to ever greater lengths to catch the oldest taxi driver available to maximize their chances of comprehension. Hotel managers are complaining that they can no longer find enough English-speakers to act as receptionists. Departing tourists, polled at the airport, voice growing frustration at not being understood (The Economist 1/20/2001).

V. Evaluating Arguments

Q: What does it mean for an argument to be good? What are the different ways in which arguments can be good? Good arguments:

1. Are persuasive.
2. Have premises that provide good evidence for the conclusion.
3. Contain premises that are true.
4. Reach a true conclusion.
5. Provide the audience good reasons for accepting the conclusion.

The focus of logic is primarily about one type of goodness: The logical relationship between premises and conclusion.

An argument is good in this sense if the premises provide good evidence for the conclusion. But what does it mean for premises to provide good evidence? We need some new concepts to capture this idea of premises providing good logical support. In order to do so, we will first need to distinguish between two types of argument.

VI. Two Types of Arguments

The two main types of arguments are called **deductive** and **inductive** arguments. We differentiate them in terms of the type of support that the premises are meant to provide for the conclusion.

Q: What function do the premises play?

Deductive Arguments are arguments in which the premises are meant to provide conclusive logical support for the conclusion.

Examples:

#1:

1. All humans are mortal
 2. Socrates is a human.
- ∴ Therefore, Socrates is mortal.

#2:

1. No student in this class will fail.
2. Mary is a student in this class.
- ∴ Therefore, Mary will not fail.

#3:

1. A intersects lines B and C.
2. Lines A and B form a 90-degree angle
3. Lines A and C form a 90-degree angle.
- ∴ B and C are parallel lines.

Inductive arguments are, by their very nature, risky arguments. Arguments in which premises provide *probable support* for the conclusion.

Statistical Examples:

#1:

1. Ten percent of all customers in this restaurant order soda.
2. John is a customer.
- ∴ John will not order Soda..

#2:

1. Some students work on campus.
2. Bill is a student.
- ∴ Bill works on campus.

#3:

1. Vegas has the Carolina Panthers as a six-point favorite for the super bowl.
- ∴ Carolina will win the Super Bowl.

VII. Good Deductive Arguments

The First Type of Goodness: Premises play their function – they provide conclusive logical support.

Deductive and inductive arguments have different aims. Deductive argument attempt to provide conclusive support or reasons; inductive argument attempt to provide probable reasons or support. So we must evaluate these two types of arguments.

Deductive arguments attempt to be valid.

An argument is **valid** if it is impossible for the premises to be true and the conclusion false.

To put validity in another way: if the premises are true, then the conclusion must be true.

It is very important to note that validity has nothing to do with whether or not the premises are, in fact, true and whether or not the conclusion is in fact true; it merely has to do with a certain conditional claim. If the premises are true, then the conclusion must be true.

Q: What does this mean?

- The validity of an argument does not depend upon the actual world. Rather, it depends upon the world described by the premises.
- First, consider the world described by the premises. In this world, is it logically possible for the conclusion to be false? That is, can you even imagine a world in which the conclusion is false?

Reflection Questions:

1. Can you have a valid argument with false premises?
 - If you cannot, then why not?
 - If you can, then provide an example of a valid argument.
2. Can you have a valid argument with a false conclusion?
 - If you cannot, then why not?
 - If you can, then provide an example of a valid argument.

You should convince yourself that validity is not just about the actual truth or falsity of the premises and conclusion. Rather, validity only has to do with a certain logical relationship between the truth of the premise and the truth of the conclusion. So the only possible combination that is ruled out by a valid argument is a set of true premises and false conclusion.

Let's go back to example #1. Here are the premises:

1. All humans are mortal.
2. Socrates is a human.

If both of these premises are true, then every human that we find must be a mortal. And this means, that it must be the case that if Socrates is a human, that Socrates is mortal.

Reflection Questions about Invalid Arguments:

1. Can you have an invalid argument with a true premise?
2. Can you have an invalid argument with true premises and a true conclusion?

The second type of goodness for deductive arguments: The premises provide us the right reasons to accept the conclusion.

Soundness Versus Validity:

An argument is sound if it is valid AND the premises are true.

Our original argument is a sound one:

1. All humans are mortal.
2. Socrates is a human.
- ∴ Socrates is mortal.

Question: *Can a sound argument have a false conclusion?*

VIII. From Deductive Arguments to Inductive Arguments

Question: *What happens if we mix around the premises and conclusion?*

1. All humans are mortal.
2. Socrates is mortal.
- ∴ Socrates is a human.

1. Socrates is mortal
2. Socrates is a human.
- ∴ All humans are mortal.

Are these valid deductive arguments?

NO, but they are common inductive arguments.

Other examples:

Suppose that there are two opaque glass jars with different color marbles in them.

1. All the marbles in jar #1 are blue.
2. This marble is blue.
- ∴ This marble came from jar #1.

1. This marble came from jar #2.
2. This marble is red.
- ∴ All the marbles in jar #2 are red.

While this is a very risky argument, what if we drew 100 marbles from jar #2 and found that they were all red? Would this affect the second argument's validity?

IX. Inductive Arguments:

The aim of an inductive argument is different from the aim of deductive argument because the type of reasons we are trying to provide are different. Therefore, the function of the premises is different in deductive and inductive arguments. And again, we can split up goodness into two types when considering inductive arguments:

1. The premises provide the right logical support.
2. The premises provide the right type of reason.

Logical Support:

Remember that for inductive arguments, the premises are intended to provide probable support for the conclusion. Thus, we

shall begin by discussing a fairly rough, coarse-grained way of talking about probable support by introducing the notions of strong and weak inductive arguments.

A strong inductive argument:

1. The vast majority of Europeans speak at least two languages.
2. Sam is a European.

∴ Sam speaks two languages.

Weak inductive argument:

1. This quarter is a fair coin.

∴ Therefore, the next coin flip will land heads.

Weak inductive argument:

1. At least one dog in this town has rabies.
2. Fido is a dog that lives in this town.

∴ Fido has rabies.

The Right Type of Reasons. As we noted above, the right type of reasons are true statements. So what happens when we get an inductive argument that is good in the first sense (right type of logical support) and good in the second sense (the right type of reasons)? Corresponding to the notion of soundness for deductive arguments, we call inductive arguments that are good in both senses cogent arguments.

A **cogent** inductive argument: a strong inductive argument with true premises

Questions:

- With which of the following types of premises and conclusions can you have a strong inductive argument?
- With which of the following types of premises and conclusions can you have a cogent inductive argument?

Premise	Conclusion
True	True
True	False
False	True
False	False

X. Steps for Evaluating Arguments:

1. Read a passage and assess whether or not it contains an argument.
2. If it does contain an argument, then identify the conclusion and premises.
3. Is this a valid deductive argument?
 - If yes, then assess it for soundness.
 - If not, then treat it as an inductive argument (step 3).

4. Is the inductive argument strong or weak?
 - If the inductive argument is strong, then is it cogent?

XI. Evaluating Real–World Arguments

An important part of evaluating arguments is not to represent the arguments of others in a deliberately weak way.

For example, suppose that I state the following:

All humans are mortal, so Socrates is mortal.

Is this valid? Not as it stands. But clearly, I believe that Socrates is a human being. Or I thought that was assumed in the conversation. That premise was clearly an implicit one.

So one of the things we can do in the evaluation of argument is to take an argument as it is stated, and represent it in a way such that it is a valid deductive argument or a strong inductive one. In doing so, we are making explicit what one would have to assume to provide a good argument (in the sense that the premises provide good – conclusive or probable – reason to accept the conclusion).

Example 1:

The teacher's policy on extra credit was unfair because Sally was the only person to have a chance at receiving extra credit.

1. Sally was the only person to have a chance at receiving extra credit.
2. The teacher's policy on extra credit is fair only if everyone gets a chance to receive extra credit.

Therefore, the teacher's policy on extra credit was unfair.

Valid argument

Example 2:

Sally didn't train very hard so she didn't win the race.

Valid:

1. Sally didn't train very hard.
2. If you don't train hard, you won't win the race.

Therefore, Sally didn't win the race.

Strong (not valid):

1. Sally didn't train very hard.
2. If you won the race, you trained hard.

Therefore, Sally didn't win the race.

Strong:

1. Sally didn't train very hard.
2. Those who don't train hard are likely not to win.

Therefore, Sally didn't win.

Example 3:

Ordinary workers receive worker's compensation benefits if they suffer an on-the-job injury. However, universities have no obligations to pay similar compensation to student athletes if they are hurt while playing sports. So, universities are not doing what they should.

1. Ordinary workers receive worker's compensation benefits if they suffer an on-the-job injury that prevents them working.
2. Student athletes are just like ordinary workers except that their job is to play sports.

3. So if student athletes are injured while playing sports, they should also be provided worker's compensation benefits.
4. Universities have no obligations to provide injured student athletes compensation.

Therefore, universities are not doing what they should.

Deductively valid argument

Example 4:

If Obama couldn't implement a single-payer healthcare system in his first term as president, then the next president will not be able to implement a single-payer healthcare system.

1. Obama couldn't implement a single-payer healthcare system.
2. In Obama's first term as president, both the House and Senate were under Democratic control.
3. The next president will either be dealing with the Republican-controlled house and senate or at best, a split legislature.
4. Obama's first term as president will be much easier than the next president's term in terms of passing legislation.

Therefore, the next president will not be able to implement a single-payer healthcare system.

Strong inductive argument

Example 5:

Sam is weaker than John. Sam is slower than John. So Sam's time on the obstacle will be slower than John's.

1. Sam is weaker than John.
2. Sam is slower than John.
3. A person's strength and speed inversely correlate with their time on the obstacle course.

Therefore, Sam's time will be slower than John's.

XII. Diagramming Arguments

All the arguments we've dealt with – except for the last two – have been fairly simple in that the premises always provided direct support for the conclusion. But in many arguments, such as the last one, there are often arguments within arguments.

Obama example:

1. Obama couldn't implement a single-payer healthcare system.
2. In Obama's first term as president, both the House and Senate were under Democratic control.
3. The next president will either be dealing with the Republican controlled house and senate or at best, a split legislature.
4. Obama's first term as president will be much easier than the next president's term in terms of passing legislation.

∴ The next president will not be able to implement a single-payer healthcare system.

It's clear that premises #2 and #3 are used in support of #4. And #1 in combination with #4 provides support for the conclusion.

When we diagram arguments, the aim is to represent the logical relationships between premises and conclusion. More specifically, we want to identify what each premise supports and how.

2 + 3



4 + 1



Ⓒ


This represents that $2+3$ together
provide support for 4
This represents that $4+1$ together
provide support for 5

When we say that $2+3$ together or $4+1$ together support some statement, we mean that the logical support of these statements are dependent upon each other. Without the other, these statements would not provide evidence for the conclusion. In order to identify when statements are dependent upon one another, we simply underline the set that are logically dependent upon one another for their evidential support. Every argument has a single conclusion, which the premises support; therefore, every argument diagram should point to the conclusion (c).

Sam Example:

1. Sam is weaker than John.
2. Sam is less flexible than John.
3. A person's strength and flexibility inversely correlate with their time on the obstacle course.

∴ Therefore, Sam's time will be slower than John's.




In some cases, different sets of premises provide evidence for the

conclusion independently of one another. In the argument above, there are two logically independent arguments for the conclusion that Sam's time will be slower than John's. That Sam is weaker than John and that being weaker correlates with a slower time provide evidence for the conclusion that Sam will be slower than John. Completely independent of this argument is the fact that Sam is less flexible and that being less flexible corresponds with a slower time. The diagram above represent these logical relations by showing that #1 and #3 dependently provide support for #4. Independent of that argument, #2 and #3 also dependently provide support for #4. Therefore, there are two logically independent sets of premises that provide support for the conclusion.

Try diagramming the following argument for yourself. The structure of the argument has been provided below:

1. All humans are mortal
2. Socrates is human
3. So Socrates is mortal.
4. If you feed a mortal person poison, he will die.

∴ Therefore, Socrates has been fed poison, so he will die.



Critical Thinking by Brian Kim is licensed under a Creative Commons Attribution 4.0 International License, except where otherwise noted.

Notes

1. This section is taken from <http://philosophy.hku.hk/think/> and is in use under creative commons license. Some modifications have been made to the original content.

3. Fallacies

I. What Are Fallacies?¹

Fallacies are mistakes of reasoning, as opposed to making mistakes that are of a factual nature. If I counted twenty people in the room when there were in fact twenty-one, then I made a factual mistake. On the other hand, if I believe that there are round squares I believe something that is contradictory. A belief in “round squares” is a mistake of reasoning and contains a fallacy because, if my reasoning were good, I would not believe something that is logically inconsistent with reality.

In some discussions, a fallacy is taken to be an undesirable kind of argument or inference. In our view, this definition of fallacy is rather narrow, since we might want to count certain mistakes of reasoning as fallacious even though they are not presented as arguments. For example, making a contradictory claim seems to be a case of fallacy, but a single claim is not an argument. Similarly, putting forward a question with an inappropriate presupposition might also be regarded as a fallacy, but a question is also not an argument. In both of these situations though, the person is making a mistake of reasoning since they are doing something that goes against one or more principles of correct reasoning. This is why we would like to define fallacies more broadly as *violations of the principles of critical thinking*, whether or not the mistakes take the form of an argument.

The study of fallacies is an application of the principles of critical thinking. Being familiar with typical fallacies can help us avoid them and help explain other people's mistakes.

There are different ways of classifying fallacies. Broadly speaking, we might divide fallacies into four kinds:

- **Fallacies of inconsistency:** cases where something inconsistent or self-defeating has been proposed or accepted.
- **Fallacies of relevance:** cases where irrelevant reasons are being invoked or relevant reasons being ignored.
- **Fallacies of insufficiency:** cases where the evidence supporting a conclusion is insufficient or weak.
- **Fallacies of inappropriate presumption:** cases where we have an assumption or a question presupposing something that is not reasonable to accept in the relevant conversational context.

II. Fallacies of Inconsistency

Fallacies of inconsistency are cases where something inconsistent, self-contradictory or self-defeating is presented.

I. Inconsistency

Here are some examples:

- “One thing that we know for certain is that nothing is ever true or false.” – If there is something we know for certain, then there is at least one truth that we know. So it can’t be the case that nothing is true or false.
- “Morality is relative and is just a matter of opinion, and so it is always wrong to impose our opinions on other people.” – But if morality is relative, it is also a relative matter whether we should impose our opinions on other people. If we should not

- do that, there is at least one thing that is objectively wrong.
- “All general claims have exceptions.” – This claim itself is a general claim, and so if it is to be regarded as true we must presuppose that there is an exception to it, which would imply that there exists at least one general claim *that does not* have an exception. So the claim itself is inconsistent.

2. Self-Defeating Claims

A self-defeating statement is a statement that, strictly speaking, is not logically inconsistent but is instead obviously false. Consider these examples:

- Very young children are fond of saying “I am not here” when they are playing hide-and-seek. The statement itself is not logically consistent, since it is not logically possible for the child not to be where she is. What is impossible is to *utter the sentence as a true sentence* (unless it is used for example in a telephone recorded message.)
- Someone who says, “I cannot speak any English.”
- Here is an actual example: A TV program in Hong Kong was critical of the Government. When the Hong Kong Chief Executive Mr. Tung was asked about it, he replied, “I shall not comment on such distasteful programs.” Mr. Tung’s remark was not logically inconsistent, because what it describes is a possible state of affairs. But it is nonetheless self-defeating because calling the program “distasteful” is to pass a comment!

III. Fallacies of Relevance

1. Taking irrelevant considerations into account

This includes defending a conclusion by appealing to irrelevant reasons, e.g., inappropriate appeal to pity, popular opinion, tradition, authority, etc. An example would be when a student failed a course and asked the teacher to give him a pass instead, because “his parents will be upset.” Since grades should be given on the basis of performance, the reason being given is quite irrelevant.

Similarly, suppose someone criticizes the Democratic Party’s call for direct elections in Hong Kong as follows: “These arguments supporting direct elections have no merit because they are advanced by Democrats who naturally stand to gain from it.” This is again fallacious because whether the person advancing the argument has something to gain from direct elections is a completely different issue from whether there ought to be direct elections.

2. Failing to Take Relevant Considerations into Account

For example, it is not unusual for us to ignore or downplay criticisms because we do not like them, even when those criticisms are justified. Or sometimes we might be tempted to make a snap decision, believing knee-jerk reactions are the best when, in fact, we should be investigating the situation more carefully and doing more research.

Of course, if we fail to consider a relevant fact simply because we are ignorant of it, then this lack of knowledge does not constitute a fallacy.

IV. Fallacies of Insufficiency

Fallacies of insufficiency are cases where insufficient evidence is provided in support of a claim. Most common fallacies fall within this category. Here are a few popular types:

1. Limited Sampling

- Momofuku Ando, the inventor of instant noodles, died at the age of 96. He said he ate instant noodles every day. So instant noodles cannot be bad for your health.
- A black cat crossed my path this morning, and I got into a traffic accident this afternoon. Black cats are really unlucky.

In both cases the observations are relevant to the conclusion, but a lot more data is needed to support the conclusion, e.g., studies show that many other people who eat instant noodles live longer, and those who encounter black cats are more likely to suffer from accidents.

2. Appeal to Ignorance

- We have no evidence showing that he is innocent. So he must be guilty.

If someone is guilty, it would indeed be hard to find evidence showing that he is innocent. But perhaps there is no evidence to point either way, so a lack of evidence is not enough to prove guilt.

3. Naturalistic Fallacy

- Many children enjoy playing video games, so we should not stop them from playing.

Many naturalistic fallacies are examples of fallacy of insufficiency. Empirical facts by themselves are not sufficient for normative conclusions, even if they are relevant.

There are many other kinds of fallacy of insufficiency. See if you can identify some of them.

V. Fallacies of Inappropriate Presumption

Fallacies of inappropriate presumption are cases where we have explicitly or implicitly made an assumption that is not reasonable to accept in the relevant context. Some examples include:

- Many people like to ask whether human nature is good or evil. This presupposes that there is such a thing as human nature and that it must be either good or bad. But why should these assumptions be accepted, and are they the only options available? What if human nature is neither good nor bad? Or what if good or bad nature applies only to individual human beings?
- Consider the question “Have you stopped being an idiot?” Whether you answer “yes” or “no,” you admit that you are, or have been, an idiot. Presumably you do not want to make any such admission. We can point out that this question has a false assumption.
- “Same-sex marriage should not be allowed because by definition a marriage should be between a man and a woman.” This argument assumes that only a heterosexual conception of

marriage is correct. But this begs the question against those who defend same-sex marriages and is not an appropriate assumption to make when debating this issue.

VI. List of Common Fallacies

ad hominem

A theory is discarded not because of any evidence against it or lack of evidence for it, but because of the person who argues for it. Example:

A: The Government should enact minimum-wage legislation so that workers are not exploited.

B: Nonsense. You say that only because you cannot find a good job.

ad ignorantiam (appeal to ignorance)

The truth of a claim is established only on the basis of lack of evidence against it. A simple obvious example of such fallacy is to argue that unicorns exist because there is no evidence against their existence. At first sight it seems that many theories that we describe as “scientific” involve such a fallacy. For example, the first law of thermodynamics holds because so far there has not been any negative instance that would serve as evidence against it. But notice, as in cases like this, there is evidence for the law, namely positive instances. Notice also that this fallacy does not apply to situations where there are only two rival claims and one has already been falsified. In situations such as this, we may justly establish the truth of the other even if we cannot find evidence for or against it.

ad misericordiam (appeal to pity)

In offering an argument, pity is appealed to. Usually this

happens when people argue for special treatment on the basis of their need, e.g., a student argues that the teacher should let them pass the examination because they need it in order to graduate. Of course, pity might be a relevant consideration in certain conditions, as in contexts involving charity.

***ad populum* (appeal to popularity)**

The truth of a claim is established only on the basis of its popularity and familiarity. This is the fallacy committed by many commercials. Surely you have heard of commercials implying that we should buy a certain product because it has made to the top of a sales rank, or because the brand is the city's "favorite."

Affirming the consequent

Inferring that P is true solely because Q is true and it is also true that if P is true, Q is true.

The problem with this type of reasoning is that it ignores the possibility that there are other conditions apart from P that might lead to Q. For example, if there is a traffic jam, a colleague may be late for work. But if we argue from his being late to there being a traffic jam, we are guilty of this fallacy – the colleague may be late due to a faulty alarm clock.

Of course, if we have evidence showing that P is the *only* or *most likely condition* that leads to Q, then we can infer that P is likely to be true without committing a fallacy.

Begging the question (*petito principii*)

In arguing for a claim, the claim itself is already assumed in the premise. Example: "God exists because this is what the Bible says, and the Bible is reliable because it is the word of God."

Complex question or loaded question

A question is posed in such a way that a person, no matter what answer they give to the question, will inevitably commit themselves to some other claim, which should not be presupposed in the context in question.

A common tactic is to ask a yes-no question that tricks people into agreeing to something they never intended to say. For example, if you are asked, “Are you still as self-centered as you used to be?”, no matter whether you answer “yes” or “no,” you are bound to admit that you were self-centered in the past. Of course, the same question would not count as a fallacy if the presupposition of the question were indeed accepted in the conversational context, i.e., that the person being asked the question had been verifiably self-centered in the past.

Composition (opposite of division)

The whole is assumed to have the same properties as its parts. Anne might be humorous and fun-loving and an excellent person to invite to the party. The same might be true of Ben, Chris and David, considered individually. But it does not follow that it will be a good idea to invite all of them to the party. Perhaps they hate each other and the party will be ruined.

Denying the antecedent

Inferring that Q is false just because if P is true, Q is also true, but P is false.

This fallacy is similar to the fallacy of affirming the consequent. Again the problem is that some alternative explanation or cause might be overlooked. Although P is false, some other condition might be sufficient to make Q true.

Example: If there is a traffic jam, a colleague may be late for work. But it is not right to argue in the light of smooth

traffic that the colleague will not be late. Again, his alarm clock may have stopped working.

Division (opposite of composition)

The parts of a whole are assumed to have the same properties as the whole. It is possible that, on a whole, a company is very effective, while some of its departments are not. It would be inappropriate to assume they all are.

Equivocation

Putting forward an argument where a word changes meaning without having it pointed out. For example, some philosophers argue that all acts are selfish. Even if you strive to serve others, you are still acting selfishly because your act is just to satisfy your desire to serve others. But surely the word “selfish” has different meanings in the premise and the conclusion – when we say a person is selfish we usually mean that he does not strive to serve others. To say that a person is selfish because he is doing something he wants, even when what he wants is to help others, is to use the term “selfish” with a different meaning.

False dilemma

Presenting a limited set of alternatives when there are others that are worth considering in the context. Example: “Every person is either my enemy or my friend. If they are my enemy, I should hate them. If they’re my friend, I should love them. So I should either love them or hate them.” Obviously, the conclusion is too extreme because most people are neither your enemy nor your friend.

Gambler’s fallacy

Assumption is made to take some independent statistics as dependent. The untrained mind tends to think that, for example, if a fair coin is tossed five times and the results are

all heads, then the next toss will more likely be a tail. It will not be, however. If the coin is fair, the result for each toss is completely independent of the others. Notice the fallacy hinges on the fact that the final result is not known. Had the final result been known already, the statistics would have been dependent.

Genetic fallacy

Thinking that because X derives from Y, and because Y has a certain property, that X must also possess that same property. Example: “His father is a criminal, so he must also be up to no good.”

Non sequitur

A conclusion is drawn that does not follow from the premise. This is not a specific fallacy but a very general term for a bad argument. So a lot of the examples above and below can be said to be non sequitur.

Post hoc, ergo propter hoc (literally, “after this, therefore because of this”)

Inferring that X must be the cause of Y just because X is followed by Y.

For example, having visited a graveyard, I fell ill and infer that graveyards are spooky places that cause illnesses. Of course, this inference is not warranted since this might just be a coincidence. However, a lot of superstitious beliefs commit this fallacy.

Red herring

Within an argument some irrelevant issue is raised that diverts attention from the main subject. The function of the red herring is sometimes to help express a strong, biased opinion. The red herring (the irrelevant issue) serves to

increase the force of the argument in a very misleading manner.

For example, in a debate as to whether God exists, someone might argue that believing in God gives peace and meaning to many people's lives. This would be an example of a red herring since whether religions can have a positive effect on people is irrelevant to the question of the existence of God. The positive psychological effect of a belief is not a reason for thinking that the belief is true.

Slippery slope

Arguing that if an opponent were to accept some claim C_1 , then they have to accept some other closely related claim C_2 , which in turn commits the opponent to a still further claim C_3 , eventually leading to the conclusion that the opponent is committed to something absurd or obviously unacceptable.

This style of argumentation constitutes a fallacy only when it is inappropriate to think if one were to accept the initial claim, one must accept all the other claims.

An example: "The government should not prohibit drugs. Otherwise the government should also ban alcohol or cigarettes. And then fatty food and junk food would have to be regulated too. The next thing you know, the government would force us to brush our teeth and do exercises every day."

Straw man

Attacking an opponent while falsely attributing to them an implausible position that is easily defeated.

Example: When many people argue for more democracy in Hong Kong, a typical "straw man" reply is to say that more democracy is not warranted because it is wrong to believe that democracy is the solution to all of Hong Kong's problems. But those who support more democracy in Hong Kong never suggest that democracy can solve *all* problems

(e.g., pollution), and those who support more democracy in Hong Kong might even agree that *blindly* accepting anything is rarely the correct course of action, whether it is democracy or not. These criticisms attack implausible “straw man” positions and do not address the real arguments for democracy.

Suppressed evidence

Where there is contradicting evidence, only confirming evidence is presented.

VII. Exercises

Identify any fallacy in each of these passages. If no fallacy is committed, select “no fallacy involved.”

1. Mr. Lee’s views on Japanese culture are wrong. This is because his parents were killed by the Japanese army during World War II and that made him anti-Japanese all his life.

2. Every ingredient of this soup is tasty. So this must be a very tasty soup.

3. Smoking causes cancer because my father was a smoker and he died of lung cancer.

4. Professor Lewis, the world authority on logic, claims that all wives cook for their husbands. But the fact is that his own wife does not cook for him. Therefore, his claim is false.

5. If Catholicism is right, then no women should be allowed to be priests. But Catholicism is wrong. Therefore, some women should be allowed to be priests.

6. God does not exist because every argument for the existence of God has been shown to be unsound.

7. The last three times I have had a cold I took large doses of vitamin C. On each occasion, the cold cleared up within a few days. So vitamin C helped me recover from colds.

8. The union's case for more funding for higher education can be ignored because it is put forward by the very people – university staff – who would benefit from the increased money.

9. Children become able to solve complex problems and think of physical objects objectively at the same time that they learn language. Therefore, these abilities are caused by learning a language.

10. If cheap things are no good then this cheap watch is no good. But this watch is actually quite good. So some good things are cheap.

1 This chapter is taken from <http://philosophy.hku.hk/think/> and is in use under the creative commons license. Some modifications have been made to the original content.



Critical Thinking by Brian Kim is licensed under a Creative Commons Attribution 4.0 International License, except where otherwise noted.

4. Sentential Logic

This chapter introduces a logical language called SL. It is a version of *sentential logic*, because the basic units of the language will represent entire sentences.

I. Sentence letters

In SL, capital letters are used to represent basic sentences. Considered only as a symbol of SL, the letter A could mean any sentence. So when translating from English into SL, it is important to provide a *symbolization key*. The key provides an English language sentence for each sentence letter used in the symbolization.

For example, consider this argument:

There is an apple on the desk.

If there is an apple on the desk, then Jenny made it to class.

∴ Jenny made it to class.

This is obviously a valid argument in English. In symbolizing it, we want to preserve the structure of the argument that makes it valid. What happens if we replace each sentence with a letter? Our symbolization key would look like this:

A: There is an apple on the desk.

B: If there is an apple on the desk, then Jenny made it to class.

C: Jenny made it to class.

We would then symbolize the argument in this way:

$$\begin{array}{c} A \\ B \\ \therefore C \end{array}$$

There is no necessary connection between some sentence *A*, which could be any sentence, and some other sentences *B* and *C*, which could be any sentences. The structure of the argument has been completely lost in this translation.

The important thing about the argument is that the second premise is not merely *any* sentence, logically divorced from the other sentences in the argument. The second premise contains the first premise and the conclusion *as parts*. Our symbolization key for the argument only needs to include meanings for *A* and *C*, and we can build the second premise from those pieces. So we symbolize the argument this way:

$$\begin{array}{c} A \\ \text{If } A, \text{ then } C. \\ \therefore C \end{array}$$

This preserves the structure of the argument that makes it valid, but it still makes use of the English expression ‘If. . . then. . .’ Although we ultimately want to replace all of the English expressions with logical notation, this is a good start.

The sentences that can be symbolized with sentence letters are called *atomic sentences*, because they are the basic building blocks out of which more complex sentences can be built. Whatever logical structure a sentence might have is lost when it is translated as an atomic sentence. From the point of view of SL, the sentence is just a

letter. It can be used to build more complex sentences, but it cannot be taken apart.

There are only twenty-six letters of the alphabet, but there is no logical limit to the number of atomic sentences. We can use the same letter to symbolize different atomic sentences by adding a subscript, a small number written after the letter. We could have a symbolization key that looks like this:

- A₁: The apple is under the armoire.
- A₂: Arguments in SL always contain atomic sentences.
- A₃: Adam Ant is taking an airplane from Anchorage to Albany.
- ⋮
- A₂₉₄: Alliteration angers otherwise affable astronauts.

Keep in mind that each of these is a different sentence letter. When there are subscripts in the symbolization key, it is important to keep track of them.

II. Connectives

Logical connectives are used to build complex sentences from atomic components. There are five logical connectives in SL. This table summarizes them, and they are explained below.

symbol	what it is called	what it means
\neg	negation	‘It is not the case that...’
$\&$	conjunction	‘Both...and...’
\vee	disjunction	‘Either...or...’
\rightarrow	conditional	‘If...then...’
\leftrightarrow	biconditional	‘...if and only if...’

Negation

Consider how we might symbolize these sentences:

1. Mary is in Barcelona.
2. Mary is not in Barcelona.
3. Mary is somewhere besides Barcelona.

In order to symbolize sentence 1, we will need one sentence letter. We can provide a symbolization key:

B: Mary is in Barcelona.

Note that here we are giving B a different interpretation than we did in the previous section. The symbolization key only specifies what B means *in a specific context*. It is vital that we continue to use this meaning of B so long as we are talking about Mary and Barcelona. Later, when we are symbolizing different sentences, we can write a new symbolization key and use B to mean something else.

Now, sentence 1 is simply B.

Since sentence 2 is obviously related to the sentence 1, we do not want to introduce a different sentence letter. To put it partly in English, the sentence means 'Not B.' In order to symbolize this, we need a symbol for logical negation. We will use '¬.' Now we can translate 'Not B' to ¬B.

Sentence 3 is about whether or not Mary is in Barcelona, but it does not contain the word 'not.' Nevertheless, it is obviously logically equivalent to sentence 2.

They both mean: It is not the case that Mary is in Barcelona. As such, we can translate both sentence 2 and sentence 3 as $\neg B$.

A sentence can be symbolized as $\neg A$ if it can be paraphrased in English as “It is not the case that A.”

Consider these further examples:

4. The widget can be replaced if it breaks.
5. The widget is irreplaceable.
6. The widget is not irreplaceable.

If we let R mean ‘The widget is replaceable’, then sentence 4 can be translated as R .

What about sentence 5? Saying the widget is irreplaceable means that it is not the case that the widget is replaceable. So even though sentence 5 is not negative in English, we symbolize it using negation as $\neg R$.

Sentence 6 can be paraphrased as ‘It is not the case that the widget is irreplaceable.’ Using negation twice, we translate this as $\neg\neg R$. The two negations in a row each work as negations, so the sentence means ‘It is not the case that . . . it is not the case that . . . R .’ If you think about the sentence in English, it is logically equivalent to sentence 4. So when we define logical equivalence in SL, we will make sure that R and $\neg\neg R$ are logically equivalent.

More examples:

7. Elliott is happy.
8. Elliott is unhappy.

If we let H mean ‘Elliot is happy’, then we can symbolize sentence 7 as H .

However, it would be a mistake to symbolize sentence 8 as $\neg H$. If Elliott is unhappy, then he is not happy– but sentence 8 does not mean the same thing as ‘It is not the case that Elliott is happy.’ It could be that he is not happy but that he is not unhappy either. Perhaps he is somewhere between the two. In order to allow for the possibility that he is indifferent, we would need a new sentence letter to symbolize sentence 8.

For any sentence A : If A is true, then $\neg A$ is false. If $\neg A$ is true, then A is false. Using ‘T’ for true and ‘F’ for false, we can summarize this in a *characteristic truth table* for negation:

\mathcal{A}	$\neg \mathcal{A}$
T	F
F	T

We will discuss truth tables at greater length in the next chapter.

Conjunction

Consider these sentences:

- 9. Adam is athletic.
- 10. Barbara is athletic.
- 11. Adam is athletic, and Barbara is also athletic.

We will need separate sentence letters for 9 and 10, so we define this symbolization key:

A: Adam is athletic.

B: Barbara is athletic.

Sentence 9 can be symbolized as A.

Sentence 10 can be symbolized as B.

Sentence 11 can be paraphrased as 'A and B.' In order to fully symbolize this sentence, we need another symbol. We will use ' & '. We translate 'A and B' as A & B. The logical connective ' & ' is called CONJUNCTION, and A and B are each called CONJUNCTS.

Notice that we make no attempt to symbolize 'also' in sentence 11. Words like 'both' and 'also' function to draw our attention to the fact that two things are being conjoined. They are not doing any further logical work, so we do not need to represent them in SL.

Some more examples:

- 12. Barbara is athletic and energetic.
- 13. Barbara and Adam are both athletic.
- 14. Although Barbara is energetic, she is not athletic.
- 15. Barbara is athletic, but Adam is more athletic than she is.

Sentence 12 is obviously a conjunction. The sentence says two things about Barbara, so in English it is permissible to refer to Barbara only once. It might be tempting to try this when translating the argument: Since B means 'Barbara is athletic', one might paraphrase the sentences as 'B and energetic.' This would be a mistake. Once we translate part of a sentence as B, any further structure is lost. B is an atomic sentence; it is nothing more than true or false. Conversely, 'energetic' is not a sentence; on its own it is neither true nor false. We should instead paraphrase the sentence as 'B and Barbara is energetic.' Now we need to add a sentence letter to the symbolization key. Let E mean 'Barbara is energetic.' Now the sentence can be translated as B & E.

A sentence can be symbolized as A & B if it can be paraphrased in English as 'Both A, and B.' Each of the conjuncts must be a sentence.

Sentence 13 says one thing about two different subjects. It says of both Barbara and Adam that they are athletic, and in English we use the word 'athletic' only once. In translating to SL, it is important to realize that the sentence can be paraphrased as, 'Barbara is athletic, and Adam is athletic.' This translates as B & A.

Sentence 14 is a bit more complicated. The word 'although' sets up a contrast between the first part of the sentence and the second part. Nevertheless, the sentence says both that Barbara is energetic and that she is not athletic. In order to make each of the conjuncts an atomic sentence, we need to replace 'she' with 'Barbara.'

So we can paraphrase sentence 14 as, 'Both Barbara is energetic, and Barbara is not athletic.' The second conjunct contains a negation, so we paraphrase further: 'Both Barbara is energetic and it is not the case that Barbara is athletic.' This translates as $E \ \& \ \neg B$.

Sentence 15 contains a similar contrastive structure. It is irrelevant for the purpose of translating to SL, so we can paraphrase the sentence as 'Both Barbara is athletic, and Adam is more athletic than Barbara.' (Notice that we once again replace the pronoun 'she' with her name.) How should we translate the second conjunct? We already have the sentence letter A which is about Adam's being athletic and B which is about Barbara's being athletic, but neither is about one of them being more athletic than the other. We need a new sentence letter. Let R mean 'Adam is more athletic than Barbara.' Now the sentence translates as $B \ \& \ R$.

Sentences that can be paraphrased 'A, but B' or
'Although A, B' are best symbolized using conjunction: A
& B

It is important to keep in mind that the sentence letters A, B, and R are atomic sentences. Considered as symbols of SL, they have no meaning beyond being true or false. We have used them to symbolize different English language sentences that are all about people being athletic, but this similarity is completely lost when we translate to SL. No formal language can capture all the structure of the English language, but as long as this structure is not important to the argument there is nothing lost by leaving it out.

For any sentences A and B, A & B is true if and only if both A and B are true. We can summarize this in the characteristic truth table for conjunction:

\mathcal{A}	\mathcal{B}	$\mathcal{A} \vee \mathcal{B}$
T	T	T
T	F	T
F	T	T
F	F	F

Conjunction is *symmetrical* because we can swap the conjuncts without changing the truth-value of the sentence. Regardless of what A and B are, A & B is logically equivalent to B & A.

Disjunction

Consider these sentences:

16. Either Denison will play golf with me, or he will watch movies.
17. Either Denison or Ellery will play golf with me.

For these sentences we can use this symbolization key:

D: Denison will play golf with me.

E: Ellery will play golf with me.

M: Denison will watch movies.

Sentence 16 is 'Either D or M .' To fully symbolize this, we

introduce a new symbol. The sentence becomes $D \vee M$. The ' \vee ' connective is called DISJUNCTION, and D and M are called DISJUNCTS.

Sentence 17 is only slightly more complicated. There are two subjects, but the English sentence only gives the verb once. In translating, we can paraphrase it as, 'Either Denison will play golf with me, or Ellery will play golf with me.' Now it obviously translates as $D \vee E$.

A sentence can be symbolized as $A \vee B$ if it can be paraphrased in English as 'Either A, or B.' Each of the disjuncts must be a sentence.

Sometimes in English, the word 'or' excludes the possibility that both disjuncts are true. This is called an EXCLUSIVE OR. An *exclusive or* is clearly intended when it says, on a restaurant menu, 'Entrees come with either soup or salad.' You may have soup; you may have salad; but, if you want *both* soup *and* salad, then you have to pay extra.

At other times, the word 'or' allows for the possibility that both disjuncts might be true. This is probably the case with sentence 17, above. I might play with Denison, with Ellery, or with both Denison and Ellery. Sentence 17 merely says that I will play with *at least* one of them. This is called an INCLUSIVE OR.

The symbol ' \vee ' represents an *inclusive or*. So $D \vee E$ is true if D is true, if E is true, or if both D and E are true. It is false only if both D and E

are false. We can summarize this with the characteristic truth table for disjunction:

\mathcal{A}	\mathcal{B}	$\mathcal{A} \vee \mathcal{B}$
T	T	T
T	F	T
F	T	T
F	F	F

Like conjunction, disjunction is symmetrical. $\mathcal{A} \vee \mathcal{B}$ is logically equivalent to $\mathcal{B} \vee \mathcal{A}$.

These sentences are somewhat more complicated:

18. Either you will not have soup, or you will not have salad.
19. You will have neither soup nor salad.
20. You get either soup or salad, but not both.

We let S_1 mean that you get soup and S_2 mean that you get salad.

Sentence 18 can be paraphrased in this way: 'Either *it is not the case that* you get soup, or *it is not the case that* you get salad.' Translating this requires both disjunction and negation. It becomes $\neg S_1 \vee \neg S_2$.

Sentence 19 also requires negation. It can be paraphrased as, 'It is *not the case that* either that you get soup or that you get salad.' We need some way of indicating that the negation does not just negate the right or left disjunct, but rather negates the entire disjunction. In order to do this, we put parentheses around the disjunction: 'It

is not the case that $(S_1 \vee S_2)$.' This becomes simply $\neg(S_1 \vee S_2)$. Notice that the parentheses are doing important work here. The sentence $\neg S_1 \vee S_2$ would mean 'Either you will not have soup, or you will have salad.'

Sentence 20 is an *exclusive or*. We can break the sentence into two parts. The first part says that you get one or the other. We translate this as $(S_1 \vee S_2)$. The second part says that you do not get both. We can paraphrase this as, 'It is not the case both that you get soup and that you get salad.' Using both negation and conjunction, we translate this as $\neg(S_1 \& S_2)$. Now we just need to put the two parts together. As we saw above, 'but' can usually be translated as a conjunction. Sentence 20 can thus be translated as $(S_1 \vee S_2) \& \neg(S_1 \& S_2)$.

Although 'v' is an *inclusive or*, we can symbolize an *exclusive or* in SL. We just need more than one connective to do it.

Conditional

For the following sentences, let R mean 'You will cut the red wire' and B mean 'The bomb will explode.'

21. If you cut the red wire, then the bomb will explode.
22. The bomb will explode only if you cut the red wire.

Sentence 21 can be translated partially as 'If R, then B.' We will use the symbol ' \rightarrow ' to represent logical entailment. The sentence becomes $R \rightarrow B$. The connective is called a **CONDITIONAL**. The sentence on the left-hand side of the conditional (R in this example) is called the **ANTECEDENT**. The sentence on the right-hand side (B) is called the **CONSEQUENT**.

Sentence 22 is also a conditional. Since the word 'if' appears in the

second half of the sentence, it might be tempting to symbolize this in the same way as sentence 21. That would be a mistake.

The conditional $R \rightarrow B$ says that *if* R were true, *then* B would also be true. It does not say that your cutting the red wire is the *only* way that the bomb could explode. Someone else might cut the wire, or the bomb might be on a timer. The sentence $R \rightarrow B$ does not say anything about what to expect if R is false. Sentence 22 is different. It says that the only conditions under which the bomb will explode involve your having cut the red wire; i.e., if the bomb explodes, then you must have cut the wire. As such, sentence 22 should be symbolized as $B \rightarrow R$.

It is important to remember that the connective ' \rightarrow ' says only that, if the antecedent is true, then the consequent is true. It says nothing about the *causal* connection between the two events. Translating sentence 22 as $B \rightarrow R$ does not mean that the bomb exploding would somehow have caused your cutting the wire. Both sentence 21 and 22 suggest that, if you cut the red wire, your cutting the red wire would be the cause of the bomb exploding. They differ on the *logical* connection. If sentence 22 were true, then an explosion would tell us— those of us safely away from the bomb— that you had cut the red wire. Without an explosion, sentence 22 tells us nothing.

The paraphrased sentence 'A only if B' is logically
equivalent to 'If A, then B.'

'If A then B' means that if A is true then so is B. So we know that if the antecedent A is true but the consequent B is false, then the conditional 'If A then B' is false. What is the truth value of 'If A then B' under other circumstances? Suppose, for instance, that the antecedent A happened to be false. 'If A then B' would then not tell us anything about the actual truth value of the consequent B, and it is unclear what the truth value of 'If A then B' would be.

In English, the truth of conditionals often depends on what *would* be the case if the antecedent *were true*— even if, as a matter of fact, the antecedent is false. This poses a problem for translating conditionals into SL. Considered as sentences of SL, R and B in the above examples have nothing intrinsic to do with each other. In order to consider what the world would be like if R were true, we would need to analyze what R says about the world. Since R is an atomic symbol of SL, however, there is no further structure to be analyzed. When we replace a sentence with a sentence letter, we consider it merely as some atomic sentence that might be true or false.

In order to translate conditionals into SL, we will not try to capture all the subtleties of the English language ‘If. . . then. . .’ Instead, the symbol ‘ \rightarrow ’ will be a *material conditional*. This means that when A is false, the conditional $A \rightarrow B$ is automatically true, regardless of the truth value of B. If both A and B are true, then the conditional $A \rightarrow B$ is true.

In short, $A \rightarrow B$ is false if and only if A is true and B is false. We can summarize this with a characteristic truth table for the conditional.

\mathcal{A}	\mathcal{B}	$\mathcal{A} \rightarrow \mathcal{B}$
T	T	T
T	F	F
F	T	T
F	F	T

The conditional is *asymmetrical*. You cannot swap the antecedent and consequent without changing the meaning of the sentence, because $A \rightarrow B$ and $B \rightarrow A$ are not logically equivalent.

Not all sentences of the form ‘If. . . then. . .’ are conditionals. Consider this sentence:

23. If anyone wants to see me, then I will be on the porch.

If I say this, it means that I will be on the porch, regardless of whether anyone wants to see me or not— but if someone did want to see me, then they should look for me there. If we let P mean ‘I will be on the porch,’ then sentence 23 can be translated simply as P .

Biconditional

Consider these sentences:

24. The figure on the board is a triangle only if it has exactly three sides.

25. The figure on the board is a triangle if it has exactly three sides.

26. The figure on the board is a triangle if and only if it has exactly three sides.

Let T mean ‘The figure is a triangle’ and S mean ‘The figure has three sides.’

Sentence 24, for reasons discussed above, can be translated as $T \rightarrow S$.

Sentence 25 is importantly different. It can be paraphrased as, ‘If the figure has three sides, then it is a triangle.’ So it can be translated as $S \rightarrow T$.

Sentence 26 says that T is true *if and only if* S is true; we can infer S from T , and we can infer T from S . This is called a biconditional, because it entails the two conditionals $S \rightarrow T$ and $T \rightarrow S$. We will use ‘ \leftrightarrow ’ to represent the biconditional; sentence 26 can be translated as $S \leftrightarrow T$.

We could abide without a new symbol for the biconditional. Since sentence 26 means ‘ $T \rightarrow S$ and $S \rightarrow T$,’ we could translate it as $(T \rightarrow$

S) & $(S \rightarrow T)$. We would need parentheses to indicate that $(T \rightarrow S)$ and $(S \rightarrow T)$ are separate conjuncts; the expression $T \rightarrow S \& S \rightarrow T$ would be ambiguous.

Because we could always write $(A \rightarrow B) \& (B \rightarrow A)$ instead of $A \leftrightarrow B$, we do not strictly speaking *need* to introduce a new symbol for the biconditional. Nevertheless, logical languages usually have such a symbol. SL will have one, which makes it easier to translate phrases like ‘if and only if’.

$A \leftrightarrow B$ is true if and only if A and B have the same truth value. This is the characteristic truth table for the biconditional:

\mathcal{A}	\mathcal{B}	$\mathcal{A} \leftrightarrow \mathcal{B}$
T	T	T
T	F	F
F	T	F
F	F	T

III. Other symbolization

We have now introduced all of the connectives of SL. We can use them together to translate many kinds of sentences. Consider these examples of sentences that use the English-language connective ‘unless’:

27. Unless you wear a jacket, you will catch cold.
28. You will catch cold unless you wear a jacket.

Let J mean ‘You will wear a jacket’ and let D mean ‘You will catch a cold.’

We can paraphrase sentence 27 as ‘Unless J , D .’ This means that if you do not wear a jacket, then you will catch cold; with this in mind, we might translate it as $\neg J \rightarrow D$. It also means that if you do not catch a cold, then you must have worn a jacket; with this in mind, we might translate it as $\neg D \rightarrow J$.

Which of these is the correct translation of sentence 27? Both translations are correct, because the two translations are logically equivalent in SL.

Sentence 28, in English, is logically equivalent to sentence 27. It can be translated as either $\neg J \rightarrow D$ or $\neg D \rightarrow J$.

When symbolizing sentences like sentence 27 and sentence 28, it is easy to get turned around. Since the conditional is not symmetric, it would be wrong to translate either sentence as $J \rightarrow \neg D$. Fortunately, there are other logically equivalent expressions. Both sentences mean that you will wear a jacket or— if you do not wear a jacket— then you will catch a cold. So we can translate them as $J \vee D$. (You might worry that the ‘or’ here should be an *exclusive* or. However, the sentences do not exclude the possibility that you might *both* wear a jacket *and* catch a cold; jackets do not protect you from all the possible ways that you might catch a cold.)

If a sentence can be paraphrased as “Unless A , B ,” then
it can be symbolized as $A \vee B$.

IV. Sentences of SL

The sentence ‘Apples are red, or berries are blue’ is a sentence of English, and the sentence ‘ $(A \vee B)$ ’ is a sentence of SL. Although we can identify sentences of English when we encounter them, we

do not have a formal definition of ‘sentence of English’. In SL, it is possible to formally define what counts as a sentence. This is one respect in which a formal language like SL is more precise than a natural language like English.

It is important to distinguish between the logical language SL, which we are developing, and the language that we use to talk about SL. When we talk about a language, the language that we are talking about is called the object language. The language that we use to talk about the OBJECT LANGUAGE is called the METALANGUAGE.

The object language in this chapter is SL. The metalanguage is English— not conversational English, but English supplemented with some logical and mathematical vocabulary. The sentence ‘ $(A \vee B)$ ’ is a sentence in the object language, because it uses only symbols of SL. The word ‘sentence’ is not itself part of SL, however, so the sentence ‘This expression is a sentence of SL’ is not a sentence of SL. It is a sentence in the metalanguage, a sentence that we use to talk *about* SL.

In this section, we will give a formal definition for ‘sentence of SL.’ The definition itself will be given in mathematical English, the metalanguage.

Expressions

There are three kinds of symbols in SL:

sentences letters with subscripts, as needed	A, B, C,...,Z or as needed: $A_1, B_1, Z_1,$ $A_2, A_{25}, J_{375},...$
connectives	$\neg, \&, \vee, \rightarrow, \leftrightarrow$
parentheses	(,)

We define an EXPRESSION of SL as any string of symbols of SL. Take any of the symbols of SL and write them down, in any order, and you have an expression.

Well-formed formulae

Since any sequence of symbols is an expression, many expressions of SL will be gobbleddegook. A meaningful expression is called a *well-formed formula*. It is common to use the acronym *wff* ; the plural is *wffs*.

Obviously, individual sentence letters like A and G_{13} will be wffs. We can form further wffs out of these by using the various connectives. Using negation, we can get $\neg A$ and $\neg G_{13}$. Using conjunction, we can get $A \& G_{13}$, $G_{13} \& A$, $A \& A$, and $G_{13} \& G_{13}$. We could also apply negation repeatedly to get wffs $\neg\neg A$ or apply negation along with conjunction to get wffs like $\neg(A \& G_{13})$ and $(G_{13} \& G_{13})$. The possible combinations are endless, even starting with just these two sentence letters, and there are infinitely many sentence letters. So there is no point in trying to list all the wffs.

Instead, we will describe the process by which wffs can be constructed. Consider negation: Given any wff A of SL, A is a wff of SL. It is important here that A is not the sentence letter A. Rather, it is a variable that stands in for any wff at all. Notice that this variable A is not a symbol of SL, so A is not an expression of SL. Instead, it is an expression of the metalanguage that allows us to talk about infinitely many expressions of SL: all of the expressions that start with the negation symbol. Because A is part of the metalanguage, it is called a *metavariable*. We can say similar things for each of the

other connectives. For instance, if A and B are wffs of SL, then $(A \ \& \ B)$ is a wff of SL. Providing clauses like this for all of the connectives, we arrive at the following formal definition for a well-formed formula of SL:

1. Every atomic sentence is a wff.
2. If A is a wff, then $\neg A$ is a wff of SL.
3. If A and B are wffs, then $(A \ \& \ B)$ is a wff.
4. If A and B are wffs, then $(A \ \vee \ B)$ is a wff.
5. If A and B are wffs, then $(A \rightarrow B)$ is a wff.
6. If A and B are wffs, then $(A \leftrightarrow B)$ is a wff.
7. All and only wffs of SL can be generated by applications of these rules.

Notice that we cannot immediately apply this definition to see whether an arbitrary expression is a wff. Suppose we want to know whether or not $\neg\neg D$ is a wff of SL. Looking at the second clause of the definition, we know that $\neg\neg D$ is a wff if $\neg D$ is a wff. So now we need to ask whether or not $\neg D$ is a wff. Again looking at the second clause of the definition, $\neg D$ is a wff if D is. Again, D is a wff if D is a wff. Now D is a sentence letter, an atomic sentence of SL, so we know that D is a wff by the first clause of the definition. So for a compound formula like $\neg\neg D$, we must apply the definition repeatedly. Eventually we arrive at the atomic sentences from which the wff is built up.

Definitions like this are called *recursive*. Recursive definitions begin with some specifiable base elements and define ways to indefinitely compound the base elements. Just as the recursive definition allows complex sentences to be built up from simple parts, you can use it to decompose sentences into their simpler parts. To determine whether or not something meets the definition, you may have to refer back to the definition many times.

The connective that you look to first in decomposing a sentence is called the MAIN LOGICAL OPERATOR of that sentence. For

example: The main logical operator of $\neg(E \vee (F \rightarrow G))$ is negation, \neg .
The main logical operator of $(\neg E \vee (F \rightarrow G))$ is disjunction, \vee .

Sentences

Recall that a sentence is a meaningful expression that can be true or false. Since the meaningful expressions of SL are the wffs and since every wff of SL is either true or false, the definition for a sentence of SL is the same as the definition for a wff. Not every formal language will have this nice feature. In the language QL, which is developed later in the book, there are wffs which are not sentences.

The recursive structure of sentences in SL will be important when we consider the circumstances under which a particular sentence would be true or false. The sentence D is true if and only if the sentence D is false, and so on through the structure of the sentence until we arrive at the atomic components: $\neg\neg\neg D$ is true if and only if the atomic sentence D is false. We will return to this point in the next chapter.

Notational conventions

A wff like $(Q \ \& \ R)$ must be surrounded by parentheses, because we might apply the definition again to use this as part of a more complicated sentence. If we negate $(Q \ \& \ R)$, we get $\neg(Q \ \& \ R)$. If we just had $Q \ \& \ R$ without the parentheses and put a negation in front of it, we would have $\neg Q \ \& \ R$. It is most natural to read this as meaning the same thing as $(\neg Q \ \& \ R)$, something very different than $(Q \ \& \ R)$. The sentence $(Q \ \& \ R)$ means that it is not the case that both Q and R are true; Q might be false or R might be false, but the sentence does not tell us which. The sentence $(\neg Q \ \& \ R)$ means specifically that Q

is false and that R is true. As such, parentheses are crucial to the meaning of the sentence.

So, strictly speaking, Q & R without parentheses is *not* a sentence of SL. When using SL, however, we will often be able to relax the precise definition so as to make things easier for ourselves. We will do this in several ways.

First, we understand that Q & R means the same thing as (Q & R). As a matter of convention, we can leave off parentheses that occur *around the entire sentence*.

Second, it can sometimes be confusing to look at long sentences with many, nested pairs of parentheses. We adopt the convention of using square brackets '[' and ']' in place of parenthesis. There is no logical difference between $(P \vee Q)$ and $[P \vee Q]$, for example. The unwieldy sentence $((H \rightarrow I) \vee (I \rightarrow H)) \& (J \vee K)$ could be written in this way: $[(H \rightarrow I) \vee (I \rightarrow H)] \& (J \vee K)$.

Third, we will sometimes want to translate the conjunction of three or more sentences. For the sentence 'Alice, Bob, and Candice all went to the party', suppose we let A mean 'Alice went', B mean 'Bob went', and C mean 'Candice went.' The definition only allows us to form a conjunction out of two sentences, so we can translate it as $(A \& B) \& C$ or as $A \& (B \& C)$. There is no reason to distinguish between these, since the two translations are logically equivalent. There is no logical difference between the first, in which $(A \& B)$ is conjoined with C, and the second, in which A is conjoined with $(B \& C)$. So we might as well just write $A \& B \& C$. As a matter of convention, we can leave out parentheses when we conjoin three or more sentences.

Fourth, a similar situation arises with multiple disjunctions. 'Either Alice, Bob, or Candice went to the party' can be translated as $(A \vee B) \vee C$ or as $A \vee (B \vee C)$. Since these two translations are logically equivalent, we may write $A \vee B \vee C$.

These latter two conventions only apply to multiple conjunctions or multiple disjunctions. If a series of connectives includes both disjunctions and conjunctions, then the parentheses are essential;

as with $(A \& B) C$ and $A \& (B C)$. The parentheses are also required if there is a series of conditionals or biconditionals; as with $(A \rightarrow B) \rightarrow C$ and $A \leftrightarrow (B \leftrightarrow C)$.

We have adopted these four rules as *notational conventions*, not as changes to the definition of a sentence. Strictly speaking, $A B C$ is still not a sentence. Instead, it is a kind of shorthand. We write it for the sake of convenience, but we really mean the sentence $(A \vee (B \vee C))$.

If we had given a different definition for a wff, then these could count as wffs. We might have written rule 3 in this way: “If A, B, \dots, Z are wffs, then $(A \& B \& \dots \& Z)$, is a wff.” This would make it easier to translate some English sentences, but would have the cost of making our formal language more complicated. We would have to keep the complex definition in mind when we develop truth tables and a proof system. We want a logical language that is *expressively simple* and allows us to translate easily from English, but we also want a *formally simple* language. Adopting notational conventions is a compromise between these two desires.

V. Practice Exercises

* **Part A** Using the symbolization key given, translate each English-language sentence into SL.

M: Those creatures are men in suits.

C: Those creatures are chimpanzees.

G: Those creatures are gorillas.

1. Those creatures are not men in suits.
2. Those creatures are men in suits, or they are not.
3. Those creatures are either gorillas or chimpanzees.
4. Those creatures are neither gorillas nor chimpanzees.
5. If those creatures are chimpanzees, then they are neither gorillas nor men in suits.

6. Unless those creatures are men in suits, they are either chimpanzees or they are gorillas.

Part B Using the symbolization key given, translate each English-language sentence into SL.

A: Mister Ace was murdered.

B: The butler did it.

C: The cook did it.

D: The Duchess is lying.

E: Mister Edge was murdered.

F: The murder weapon was a frying pan.

1. Either Mister Ace or Mister Edge was murdered.
2. If Mister Ace was murdered, then the cook did it.
3. If Mister Edge was murdered, then the cook did not do it.
4. Either the butler did it, or the Duchess is lying.
5. The cook did it only if the Duchess is lying.
6. If the murder weapon was a frying pan, then the culprit must have been the cook.
7. If the murder weapon was not a frying pan, then the culprit was either the cook or the butler.
8. Mister Ace was murdered if and only if Mister Edge was not murdered.
9. The Duchess is lying, unless it was Mister Edge who was murdered.
10. If Mister Ace was murdered, he was done in with a frying pan.
11. Since the cook did it, the butler did not.
12. Of course the Duchess is lying!

* **Part C** Using the symbolization key given, translate each English-language sentence into SL.

E1: Ava is an electrician.
E2: Harrison is an electrician.
F1: Ava is a firefighter.
F2: Harrison is a firefighter.
S1: Ava is satisfied with her career.
S2: Harrison is satisfied with his career.

1. Ava and Harrison are both electricians.
2. If Ava is a firefighter, then she is satisfied with her career.
3. Ava is a firefighter, unless she is an electrician.
4. Harrison is an unsatisfied electrician.
5. Neither Ava nor Harrison is an electrician.
6. Both Ava and Harrison are electricians, but neither of them find it satisfying.
7. Harrison is satisfied only if he is a firefighter.
8. If Ava is not an electrician, then neither is Harrison, but if she is, then he is too.
9. Ava is satisfied with her career if and only if Harrison is not satisfied with his.
10. If Harrison is both an electrician and a firefighter, then he must be satisfied with his work.
11. It cannot be that Harrison is both an electrician and a firefighter.
12. Harrison and Ava are both firefighters if and only if neither of them is an electrician.

* **Part D** Give a symbolization key and symbolize the following sentences in SL.

1. Alice and Bob are both spies.

2. If either Alice or Bob is a spy, then the code has been broken.
3. If neither Alice nor Bob is a spy, then the code remains unbroken.
4. The German embassy will be in an uproar, unless someone has broken the code.
5. Either the code has been broken or it has not, but the German embassy will be in an uproar regardless.
6. Either Alice or Bob is a spy, but not both.

Part E Give a symbolization key and symbolize the following sentences in SL.

1. If Gregor plays first base, then the team will lose.
2. The team will lose unless there is a miracle.
3. The team will either lose or it won't, but Gregor will play first base regardless.
4. Gregor's mom will bake cookies if and only if Gregor plays first base.
5. If there is a miracle, then Gregor's mom will not bake cookies.

Part F For each argument, write a symbolization key and translate the argument as well as possible into SL.

1. If Dorothy plays the piano in the morning, then Roger wakes up cranky. Dorothy plays piano in the morning unless she is distracted. So if Roger does not wake up cranky, then Dorothy must be distracted.
2. It will either rain or snow on Tuesday. If it rains, Neville will be sad. If it snows, Neville will be cold. Therefore, Neville will either be sad or cold on Tuesday.
3. If Zoog remembered to do his chores, then things are clean but not neat. If he forgot, then things are neat but not clean.

Therefore, things are either neat or clean— but not both.

* **Part G** For each of the following: (a) Is it a wff of SL? (b) Is it a sentence of SL, allowing for notational conventions?

1. (A)
2. $J374 \vee \neg J374$
3. $\neg\neg\neg\neg F$
4. $\neg \& S$
5. $(G \& \neg G)$
6. $A \rightarrow A$
7. $(A \rightarrow (A \& \neg F)) \vee (D \leftrightarrow E)$
8. $[(Z \leftrightarrow S) \rightarrow W] \& [J \vee X]$
9. $(F \leftrightarrow \neg D \rightarrow J) \vee (C \& D)$

Part H

1. 1. Are there any wffs of SL that contain no sentence letters?
Why or why not?
2. 2. In the chapter, we symbolized an *exclusive or* using \vee , $\&$, and \neg . How could you translate an *exclusive or* using only two connectives? Is there any way to translate an *exclusive or* using only one connective?

5. Truth Tables

This chapter introduces a way of evaluating sentences and arguments of SL. Although it can be laborious, the truth table method is a purely mechanical procedure that requires no intuition or special insight.

I. Truth-functional connectives

Any non-atomic sentence of SL is composed of atomic sentences with sentential connectives. The truth-value of the compound sentence depends only on the truth-value of the atomic sentences that comprise it. In order to know the truth-value of $(D \leftrightarrow E)$, for instance, you only need to know the truth-value of D and the truth-value of E . Connectives that work in this way are called TRUTH-FUNCTIONAL.

In this chapter, we will make use of the fact that all of the logical operators in SL are truth-functional— it makes it possible to construct truth tables to determine the logical features of sentences. You should realize, however, that this is not possible for all languages. In English, it is possible to form a new sentence from any simpler sentence X by saying ‘It is possible that X .’ The truth-value of this new sentence does not depend directly on the truth-value of X . Even if X is false, perhaps in some sense X *could* have been true— then the new sentence would be true. Some formal languages, called *modal logics*, have an operator for possibility. In a modal logic, we could translate ‘It is possible that X ’ as $\Diamond X$. However, the ability to translate sentences like these come at a cost: The \Diamond operator is not truth-functional, and so modal logics are not amenable to truth tables.

\mathcal{A}	$\neg\mathcal{A}$	\mathcal{A}	\mathcal{B}	$\mathcal{A} \& \mathcal{B}$	$\mathcal{A} \vee \mathcal{B}$	$\mathcal{A} \rightarrow \mathcal{B}$	$\mathcal{A} \leftrightarrow \mathcal{B}$
1	0	1	1	1	1	1	1
1	0	1	0	0	1	0	0
0	1	0	1	0	1	1	0
0	1	0	0	0	0	1	1

Table 5.1: The characteristic truth tables for the connectives of SL.

II. Complete truth tables

The truth-value of sentences which contain only one connective are given by the characteristic truth table for that connective. In the previous chapter, we wrote the characteristic truth tables with ‘T’ for true and ‘F’ for false. It is important to note, however, that this is not about truth in any deep or cosmic sense. Poets and philosophers can argue at length about the nature and significance *truth*, but the truth functions in SL are just rules which transform input values into output values. To underscore this, in this chapter we will write ‘1’ and ‘0’ instead of ‘T’ and ‘F’. Even though we interpret ‘1’ as meaning ‘true’ and ‘0’ as meaning ‘false’, computers can be programmed to fill out truth tables in a purely mechanical way. In a machine, ‘1’ might mean that a register is switched on and ‘0’ that the register is switched off. Mathematically, they are just the two possible values that a sentence of SL can have. The truth tables for the connectives of SL, written in terms of 1s and 0s, are given in table 5.1.

The characteristic truth table for conjunction, for example, gives the truth conditions for any sentence of the form $(A \& B)$. Even if the conjuncts A and B are long, complicated sentences, the conjunction is true if and only if both A and B are true. Consider the sentence $(H \& I) \rightarrow H$. We consider all the possible combinations of true and false for H and I , which gives us four rows. We then copy the truth-values

for the sentence letters and write them underneath the letters in the sentence.

H	I	$(H \ \& \ I) \rightarrow H$			
1	1	1	1	1	
1	0	1	0	1	
0	1	0	1	0	
0	0	0	0	0	

Now consider the subsentence $H \ \& \ I$. This is a conjunction $A \ \& \ B$ with H as A and with I as B . H and I are both true on the first row. Since a conjunction is true when both conjuncts are true, we write a 1 underneath the conjunction symbol. We continue for the other three rows and get this:

H	I	$(H \ \& \ I) \rightarrow H$			
		$\mathcal{A} \ \& \ \mathcal{B}$			
1	1	1	1	1	1
1	0	1	0	0	1
0	1	0	0	1	0
0	0	0	0	0	0

The entire sentence is a conditional $A \rightarrow B$ with $(H \ \& \ I)$ as A and with H as B . On the second row, for example, $(H \ \& \ I)$ is false and H is true. Since a conditional is true when the antecedent is false, we

write a 1 in the second row underneath the conditional symbol. We continue for the other three rows and get this:

H	I	$(H \ \& \ I) \rightarrow H$		
		\mathcal{A}	\rightarrow	\mathcal{B}
1	1	1	1	1
1	0	0	1	1
0	1	0	1	0
0	0	0	1	0

The column of 1s underneath the conditional tells us that the sentence $(H \ \& \ I) \rightarrow I$ is true regardless of the truth-values of H and I . They can be true or false in any combination, and the compound sentence still comes out true. It is crucial that we have considered all of the possible combinations. If we only had a two-line truth table, we could not be sure that the sentence was not false for some other combination of truth-values.

In this example, we have not repeated all of the entries in every successive table. When actually writing truth tables on paper, however, it is impractical to erase whole columns or rewrite the whole table for every step. Although it is more crowded, the truth table can be written in this way:

H	I	$(H \ \& \ I) \rightarrow H$				
1	1	1	1	1	1	1
1	0	1	0	0	1	1
0	1	0	0	1	1	0
0	0	0	0	0	1	0

Most of the columns underneath the sentence are only there for bookkeeping purposes. When you become more adept with truth tables, you will probably no longer need to copy over the columns for each of the sentence letters. In any case, the truth-value of the sentence on each row is just the column underneath the main logical operator of the sentence; in this case, the column underneath the conditional.

A COMPLETE TRUTH TABLE has a row for all the possible combinations of 1 and 0 for all of the sentence letters. The size of the complete truth table depends on the number of different sentence letters in the table. A sentence that contains only one sentence letter requires only two rows, as in the characteristic truth table for negation. This is true even if the same letter is repeated many times, as in the sentence $[(C \leftrightarrow C) \rightarrow C] \ \& \ \neg (C \rightarrow C)$. The complete truth table requires only two lines because there are only two possibilities: C can be true or it can be false. A single sentence letter can never be marked both 1 and 0 on

the same row. The truth table for this sentence looks like this:

C	$[(C \leftrightarrow C) \rightarrow C] \ \& \ \neg (C \rightarrow C)$									
1	1	1	1	1	1	0	0	1	1	1
0	0	1	0	0	0	0	0	0	1	0

Looking at the column underneath the main connective, we see that the sentence is false on both rows of the table; i.e., it is false regardless of whether C is true or false.

A sentence that contains two sentence letters requires four lines for a complete truth table, as in the characteristic truth tables and the table for $(H \ \& \ I) \rightarrow I$.

A sentence that contains three sentence letters requires eight lines. For example:

M	N	P	$M \ \& \ (N \vee P)$				
1	1	1	1	1	1	1	1
1	1	0	1	1	1	1	0
1	0	1	1	1	0	1	1
1	0	0	1	0	0	0	0
0	1	1	0	0	1	1	1
0	1	0	0	0	1	1	0
0	0	1	0	0	0	1	1
0	0	0	0	0	0	0	0

From this table, we know that the sentence $M \ \& \ (N \vee P)$ might be true or false, depending on the truth-values of M , N , and P .

A complete truth table for a sentence that contains four different sentence letters requires 16 lines. Five letters, 32 lines. Six letters, 64 lines. And so on. To be perfectly general: If a complete truth table has n different sentence letters, then it must have 2^n rows.

In order to fill in the columns of a complete truth table, begin with

the right-most sentence letter and alternate 1s and 0s. In the next column to the left, write two 1s, write two 0s, and repeat. For the third sentence letter, write four 1s followed by four 0s. This yields an eight line truth table like the one above.

For a 16 line truth table, the next column of sentence letters should have eight 1s followed by eight 0s. For a 32 line table, the next column would have 16 1s followed by 16 0s. And so on.

III. Using truth tables

Tautologies, contradictions, and contingent sentences

Recall that an English sentence is a tautology if it must be true as a matter of logic. With a complete truth table, we consider all of the ways that the world might be. If the sentence is true on every line of a complete truth table, then it is true as a matter of logic, regardless of what the world is like.

So a sentence is a TAUTOLOGY IN SL if the column under its main connective is 1 on every row of a complete truth table.

Conversely, a sentence is a CONTRADICTION IN SL if the column under its main connective is 0 on every row of a complete truth table.

A sentence is CONTINGENT IN SL if it is neither a tautology nor a contradiction; i.e. if it is 1 on at least one row and 0 on at least one row.

From the truth tables in the previous section, we know that $(H \& I) \rightarrow H$ is a tautology, that $[(C \leftrightarrow C) \rightarrow C] \& \neg(C \rightarrow C)$ is a contradiction, and that $M \& (N \vee P)$ is contingent.

Logical equivalence

Two sentences are logically equivalent in English if they have the same truth value as a matter of logic. Once again, truth tables allow us to define an analogous concept for SL: Two sentences are LOGICALLY EQUIVALENT IN SL if they have the same truth-value on every row of a complete truth table.

Consider the sentences $\neg(A \vee B)$ and $\neg A \ \& \ \neg B$. Are they logically equivalent? To find out, we construct a truth table.

A	B	$\neg(A \vee B)$	$\neg A \ \& \ \neg B$
1	1	0	0
1	0	0	0
0	1	0	0
0	0	1	1

Look at the columns for the main connectives; negation for the first sentence, conjunction for the second. On the first three rows, both are 0. On the final row, both are 1. Since they match on every row, the two sentences are logically equivalent.

Consistency

A set of sentences in English is consistent if it is logically possible for them all to be true at once. A set of sentences is LOGICALLY CONSISTENT IN SL if there is at least one line of a complete truth table on which all of the sentences are true. It is INCONSISTENT otherwise.

Validity

An argument in English is valid if it is logically impossible for the premises to be true and for the conclusion to be false at the same time. An argument is VALID IN SL if there is no row of a complete truth table on which the premises are all 1 and the conclusion is 0; an argument is INVALID IN SL if there is such a row.

Consider this argument:

$$\neg L \rightarrow (J \vee L)$$

$$\neg L$$

$$\therefore J$$

Is it valid? To find out, we construct a truth table.

J	L	$\neg L \rightarrow (J \vee L)$						$\neg L$	J
1	1	0	1	1	1	1	1	0	1
1	0	1	0	1	1	1	0	1	1
0	1	0	1	1	0	1	1	0	0
0	0	1	0	0	0	0	0	1	0

Yes, the argument is valid. The only row on which both the

premises are 1 is the second row, and on that row the conclusion is also 1.

IV. Partial truth tables

In order to show that a sentence is a tautology, we need to show that it is 1 on every row. So we need a complete truth table. To show that a sentence is *not* a tautology, however, we only need one line: a line on which the sentence is 0. Therefore, in order to show that something is not a tautology, it is enough to provide a one-line *partial truth table*— regardless of how many sentence letters the sentence might have in it.

Consider, for example, the sentence $(U \ \& \ T) \rightarrow (S \ \& \ W)$. We want to show that it is *not* a tautology by providing a partial truth table. We fill in 0 for the entire sentence. The main connective of the sentence is a conditional. In order for the conditional to be false, the antecedent must be true (1) and the consequent must be false (0). So we fill these in on the table:

S	T	U	W	$(U \ \& \ T) \rightarrow (S \ \& \ W)$
				1 0 0

In order for the $(U \ \& \ T)$ to be true, both U and T must be true.

S	T	U	W	$(U \& T) \rightarrow (S \& W)$				
	1	1		1	1	1	0	0

Now we just need to make $(S \& W)$ false. To do this, we need to make at least one of S and W false. We can make both S and W false if we want. All that matters is that the whole sentence turns out false on this line. Making an arbitrary decision, we finish the table in this way:

S	T	U	W	$(U \& T) \rightarrow (S \& W)$				
0	1	1	0	1	1	1	0	0 0 0

Showing that something is a contradiction requires a complete truth table. Showing that something is *not* a contradiction requires only a one-line partial truth table, where the sentence is true on that one line.

A sentence is contingent if it is neither a tautology nor a contradiction. So showing that a sentence is contingent requires a *two-line* partial truth table: The sentence must be true on one line and false on the other. For example, we can show that the sentence above is contingent with this truth table:

S	T	U	W	$(U \& T) \rightarrow (S \& W)$				
0	1	1	0	1	1	1	0	0 0 0
0	1	0	0	0	0	1	1	0 0 0

	YES	NO
tautology?	complete truth table	one-line partial truth table
contradiction?	complete truth table	one-line partial truth table
contingent?	two-line partial truth table	complete truth table
equivalent?	complete truth table	one-line partial truth table
consistent?	one-line partial truth table	complete truth table
valid?	complete truth table	one-line partial truth table

Table 5.2: Do you need a complete truth table or a partial truth table? It depends on what you are trying to show.

Note that there are many combinations of truth values that would have made the sentence true, so there are many ways we could have written the second line.

Showing that a sentence is *not* contingent requires providing a complete truth table, because it requires showing that the sentence is a tautology or that it is a contradiction. If you do not know whether a particular sentence is contingent, then you do not know whether you will need a complete or partial truth table. You can always start working on a complete truth table. If you complete rows that show the sentence is contingent, then you can stop. If not, then complete the truth table. Even though two carefully selected rows will show that a contingent sentence is contingent, there is nothing wrong with filling in more rows.

Showing that two sentences are logically equivalent requires providing a complete truth table. Showing that two sentences are *not* logically equivalent requires only a one-line partial truth table: Make the table so that one sentence is true and the other false.

Showing that a set of sentences is consistent requires providing one row of a truth table on which all of the sentences are true. The rest of the table is irrelevant, so a one-line partial truth table will do. Showing that a set of sentences is inconsistent, on the other hand, requires a complete truth table: You must show that on every row of the table at least one of the sentences is false.

Showing that an argument is valid requires a complete truth table. Showing that an argument is *invalid* only requires providing a one-line truth table: If you can produce a line on which the premises are all true and the conclusion is false, then the argument is invalid.

Table 5.2 summarizes when a complete truth table is required and when a partial truth table will do.

V. Practice Exercises

If you want additional practice, you can construct truth tables for any of the sentences and arguments in the exercises for the previous chapter.

* **Part A** Determine whether each sentence is a tautology, a contradiction, or a contingent sentence. Justify your answer with a complete or partial truth table where appropriate.

1. $A \rightarrow A$
2. $\neg B \& B$
3. $C \rightarrow \neg C$
4. $\neg D \vee D$
5. $(A \leftrightarrow B) \leftrightarrow \neg(A \leftrightarrow \neg B)$
6. $(A \& B) \vee (B \& A)$
7. $(A \rightarrow B) \vee (B \rightarrow A)$
8. $\neg[A \rightarrow (B \rightarrow A)]$
9. $(A \& B) \rightarrow (B \vee A)$
10. $A \leftrightarrow [A \rightarrow (B \& \neg B)]$
11. $\neg(A \vee B) \leftrightarrow (\neg A \& \neg B)$
12. $\neg(A \& B) \leftrightarrow A$
13. $[(A \& B) \& \neg(A \& B)] \& C$
14. $A \rightarrow (B \vee C)$
15. $[(A \& B) \& C] \rightarrow B$
16. $(A \& \neg A) \rightarrow (B \vee C)$
17. $\neg[(C \vee A) \vee B]$
18. $(B \& D) \leftrightarrow [A \leftrightarrow (A \vee C)]$

* **Part B** Determine whether each pair of sentences is logically equivalent. Justify your answer with a complete or partial truth table where appropriate.

1. $A, \neg A$
2. $A, A \vee A$
3. $A \rightarrow A, A \leftrightarrow A$
4. $A \vee \neg B, A \rightarrow B$
5. $A \& \neg A, \neg B \leftrightarrow B$
6. $\neg(A \& B), \neg A \vee \neg B$
7. $\neg(A \rightarrow B), \neg A \rightarrow \neg B$
8. $(A \rightarrow B), (\neg B \rightarrow \neg A)$
9. $[(A \vee B) \vee C], [A \vee (B \vee C)]$
10. $[(A \vee B) \& C], [A \vee (B \& C)]$

* **Part C** Determine whether each set of sentences is consistent or inconsistent. Justify your answer with a complete or partial truth table where appropriate.

1. $A \rightarrow A, \neg A \rightarrow \neg A, A \& A, A \vee A$
2. $A \& B, C \rightarrow \neg B, C$
3. $A \vee B, A \rightarrow C, B \rightarrow C$
4. $A \rightarrow B, B \rightarrow C, A, \neg C$
5. $B \& (C \vee A), A \rightarrow B, \neg(B \vee C)$
6. $A \vee B, B \vee C, C \rightarrow \neg A$
7. $A \leftrightarrow (B \vee C), C \rightarrow \neg A, A \rightarrow \neg B$
8. $A, B, C, \neg D, \neg E, F$

* **Part D** Determine whether each argument is valid or invalid. Justify your answer with a complete or partial truth table where appropriate.

1. $A \rightarrow A, \therefore A$
2. $A \vee [A \rightarrow (A \leftrightarrow A)], \therefore A$
3. $A \rightarrow (A \& \neg A), \therefore \neg A$
4. $A \leftrightarrow \neg(B \leftrightarrow A), \therefore A$
5. $A \vee (B \rightarrow A), \therefore \neg A \rightarrow \neg B$
6. $A \rightarrow B, B, \therefore A$
7. $A \vee B, B \vee C, \neg A, \therefore B \& C$
8. $A \vee B, B \vee C, \neg B, \therefore A \& C$
9. $(B \& A) \rightarrow C, (C \& A) \rightarrow B, \therefore (C \& B) \rightarrow A$
10. $A \leftrightarrow B, B \leftrightarrow C, \therefore A \leftrightarrow C$

* **Part E** Answer each of the questions below and justify your answer.

1. Suppose that \mathcal{A} and \mathcal{B} are logically equivalent. What can you say about $\mathcal{A} \leftrightarrow \mathcal{B}$?
2. Suppose that $(\mathcal{A} \& \mathcal{B}) \rightarrow \mathcal{C}$ is contingent. What can you say about the argument “ $\mathcal{A}, \mathcal{B}, \therefore \mathcal{C}$ ”?
3. Suppose that $\{\mathcal{A}, \mathcal{B}, \mathcal{C}\}$ is inconsistent. What can you say about $(\mathcal{A} \& \mathcal{B} \& \mathcal{C})$?
4. Suppose that \mathcal{A} is a contradiction. What can you say about the argument “ $\mathcal{A}, \mathcal{B}, \therefore \mathcal{C}$ ”?
5. Suppose that \mathcal{C} is a tautology. What can you say about the argument “ $\mathcal{A}, \mathcal{B}, \therefore \mathcal{C}$ ”?
6. Suppose that \mathcal{A} and \mathcal{B} are logically equivalent. What can you say about $(\mathcal{A} \vee \mathcal{B})$?
7. Suppose that \mathcal{A} and \mathcal{B} are *not* logically equivalent. What can you say about $(\mathcal{A} \vee \mathcal{B})$?

Part F We could leave the biconditional (\leftrightarrow) out of the language. If we did that, we could still write ‘ $A \leftrightarrow B$ ’ so as to make sentences easier to read, but that would be shorthand for $(A \rightarrow B) \& (B \rightarrow A)$.

The resulting language would be formally equivalent to SL, since $A \leftrightarrow B$ and $(A \rightarrow B) \& (B \rightarrow A)$ are logically equivalent in SL. If we valued formal simplicity over expressive richness, we could replace more of the connectives with notational conventions and still have a language equivalent to SL.

There are a number of equivalent languages with only two connectives. It would be enough to have only negation and the material conditional. Show this by writing sentences that are logically equivalent to each of the following using only parentheses, sentence letters, negation (\neg), and the material conditional (\rightarrow).

1. $A \vee B$
2. $A \& B$
3. $A \leftrightarrow B$

We could have a language that is equivalent to SL with only negation and disjunction as connectives. Show this: Using only parentheses, sentence letters, negation (\neg), and disjunction (\vee), write sentences that are logically equivalent to each of the following.

4. $A \& B$
5. $A \rightarrow B$
6. $A \leftrightarrow B$

The *Sheffer stroke* is a logical connective with the following characteristic truth table:

\mathcal{A}	\mathcal{B}	$\mathcal{A} \mathcal{B}$
1	1	0
1	0	1
0	1	1
0	0	1

7. Write a sentence using the connectives of SL that is logically equivalent to $(\mathcal{A}|\mathcal{B})$.

Every sentence written using a connective of SL can be rewritten as a logically equivalent sentence using one or more Sheffer strokes. Using only the Sheffer stroke, write sentences that are equivalent to each of the following.

8. $\neg A$
9. $(A \& B)$
10. $(A \vee B)$
11. $(A \rightarrow B)$
12. $(A \leftrightarrow B)$

6. Categorical Logic

I. Venn diagrams^I

Pictures and diagrams can be very useful in presenting information or assisting reasoning. In this module we shall focus on Venn diagram. They are used to represent classes of objects. We can also use them to evaluate the validity of certain types of arguments.

Venn diagrams are named after the British logician John Venn (1834-1923), a fellow of Gonville and Caius college at Cambridge University. He was also a philosopher and mathematician, a pioneer of logic and probability theory.


II. Basic Notation

1. A class is defined by its members

Let us start with the concept of a class. A class or a set is simply a collection of objects. These objects are called members of the set. A class is defined by its members. So for example, we might define a class C as the class of black hats. In that case, every black hat in the world is a member of C, and anything that is not a black hat is not a member of C. If something is not a member of a class, we can also say that the object is outside the class.

Note that a class can be empty. The class of men over 5 meters tall is presumably empty since nobody is that tall. The class of plane figures that are both round and square is also empty since nothing can be both round and square. A class can also be infinite,

containing an infinite number of objects. The class of even number is an example. It has infinitely many members, including 2, 4, 6, 8, and so on.




2. Classes are represented by circles

- As you can see in the diagram above, the class of black hats, C , is represented by a circle. We normally use circles to represent classes in Venn diagrams, though sometimes we also use bounded regions with different shapes, such as ovals.
- We can write the name of the class, e.g. “ C ”, or “Class C ”, next to the circle to indicate which class it is.
- The area inside the circle represents those things which are members of the class.
- The area outside the circle represents those things which are not members of the class, e.g. green hats, keys, cakes, etc.
- A Venn diagram is usually enclosed by a rectangular box that represents everything in the world.

3. Use shading to indicate an empty class

Let us now consider what shading means:



To indicate that a class is empty, we shade the circle representing that class. So the diagram above means that class A is empty.




In general, shading an area means that the class represented by the area is empty. So the second diagram above represents a situation where there isn't anything which is not a member of class A.

However, even though shading indicates emptiness, a region that


is not shaded does not necessarily indicate a non-empty class. As we shall see in the next tutorial, we use a tick to indicate existence. So in the second diagram above, the circle marked A is not shaded. This does not imply that there are things which exist which are members of A. If the area is blank, this means that we do not have any information as to whether there is anything there.

III. Everything and nothing

I. Intersecting circles



Now let us consider a slightly more complicated diagram where we have two intersecting circles. The left circle represents class A. The right one represents class B.



Let us label the different bounded regions:

- Region 1 represents objects which belong to class A but not to B.
- Region 2 represents objects which belong to both A and B.
- Region 3 represents objects which belong to B but not A.
- Region 4, the area outside the two circles, represents objects that belong to neither A nor B.


Exercise #1

So for example, suppose A is the class of apples, and B is the class of sweet things. In that case what does region 2 represent?


Exercise #2

Furthermore, which region represents the class that contains sour lemons that are not sweet?

2. Everything and nothing




Continuing with our diagram, suppose we now shade region 1. This means that the class of things which belong to A but not B is empty. Or more simply, every A is a B. (It might be useful to note that this is equivalent to saying that if anything is an A, it is also a B.) This is an important point to remember. Whenever you want to represent “every A is B”, shade the area within the A circle that is outside the B circle.




What if we shade the middle region where A and B overlaps? This

is the region representing things which are both A and B. So shading indicates that nothing is both A and B. If you think about it carefully, you will see that “Nothing is both A and B” says the same thing as “No A is a B” and “No B is an A”. Make sure that you understand why these claims are logically equivalent!

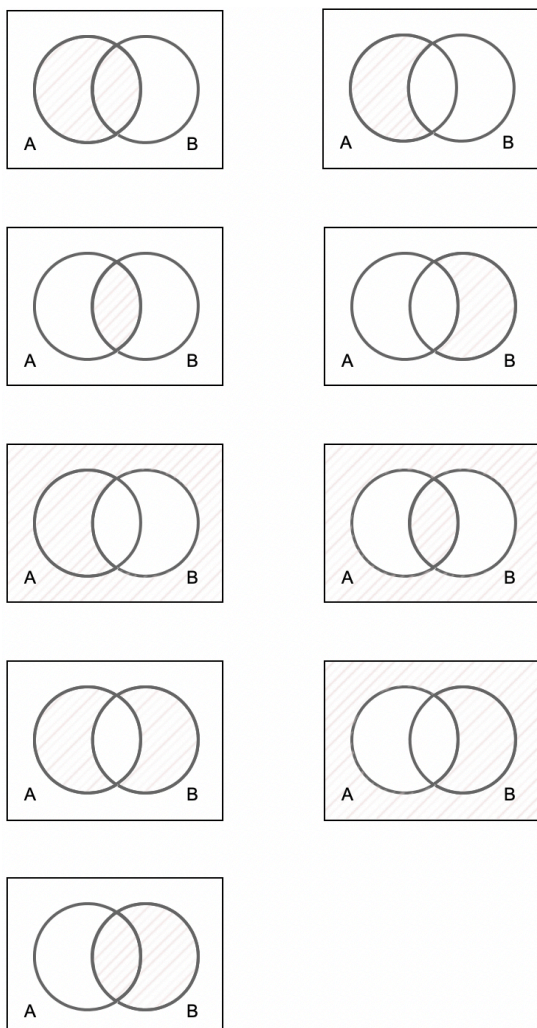


Incidentally, we could have represented the same information by using two non-overlapping circles instead.



IV. Exercises

See if you can explain what each diagram represents.




V. Three circles

So far we have been looking at Venn diagrams with two circles. We now turn to Venn diagrams with three circles. The interpretation of

these diagrams is the same as before, with each circle representing a class of objects, and the overlapping area between the circles representing the class of objects that belong to all the classes.

As you can see from the diagram below, with three circles we can have eight different regions, the eighth being the region outside the circles. The top circle represents the class of As, whereas the circles on the left and the right below it represent the class of Bs and Cs respectively. The area outside all the circles represents those objects which are not members of any of these three classes.




I. Shading

Now that you know what each of the region represents, you should know how to use shading to represent situations where “Every X is Y”, or “No X is Y”. As before, shading an area indicates that nothing exists in the class that is represented by the shaded region.

Exercise #1


Look at the sentences in the diagram below. Ask yourself which region should be shaded to represent the situation described by the sentence. Then click that sentence and check the answer.



Every A is B	No A is B
Every A is C	No A is C
Every B is A	No B is A
Every B is C	No B is C
Every C is B	No C is B
Every C is A	No C is A

VI. Existence

We have seen how to use shading to indicate that there is nothing in the class represented by the shaded region. We now see how to use ticks to indicate existence. The basic idea is that when a tick is present in a region, it indicates that there is something in the class represented by the region. So for example, in the diagram below, we have a tick outside the circles. Since the area outside the circle represents the class of things that are neither A, nor B, nor C, the diagram is saying that something exists that is neither A nor B nor C:




There are two important points to remember :


1. A tick in a region says that *there is something* in the class represented by the region. It does not say *how many* things there are in that class. There might be just one, or perhaps there are many.
2. A region without a tick *does not* represent an empty class.

Without a tick, a blank region provides no information as to whether anything exists in the class it represents. Only when a region is shaded can we say that it represents an empty class.

What about the following diagram? What does it represent?




The diagram above does NOT say “something is A”. Actually it says something more specific, namely that “something is A but is not B and not C”. If you have given the wrong answer, you might be thinking that the tick indicates that there is something in the class represented by the A circle. But here we use a tick to indicate existence in the class represented by the *smallest bounded region that encloses the tick*. In the top diagram of this page the smallest bounded area that encloses the tick is the area outside the three circles. In the diagram above, although circle A does enclose the tick, it is not the *smallest* bounded area that does that. That smallest region is the colored one in this diagram :




Now see if you can determine what these diagrams indicate.


Exercise #1




Exercise #2



Exercise #3




Exercise #4



Notice in the last diagram above, the two ticks indicate that there

are two different things. What if you just want to say “Something is C but not A”? The way to do this is to put a tick across two bounded regions, as follows:

Exercise #5




The interpretation of this diagram employs the same rule as before. What the tick indicates is that there is something in the smallest closed region (the colored area) that encloses the tick. Of course, the bigger C circle also completely encloses the tick, but it is not the *smallest* bounded region that does that. So the tick does not mean that “something is C”.


Notice that the tick does not tell us whether there is anything that is B, because it is not completely enclosed by the B circle.

See if you can explain why these diagrams represent:


Exercise #6




Exercise #7




Exercise #8



Exercise #9




So far we have used ticks to cut across only two bounded regions. But of course there are other possibilities:




What do you think this means? Applying the same rule of interpretation as before, we see that the smallest closed region that encloses the big tick would have to be the combined three regions which the tick spreads across. This combined region represents things which are either B or C (or both), but which are not A. So what the diagram says is that there is something of this kind.


So what if we just want to represent the fact that something is A? Here is one way to draw the diagram: Notice that the tick cuts across all the different regions within the A circle, and is completely enclosed by it.



We can now combine what we have learnt about ticks and shading together. Suppose we start with the information that something is both A and C. We therefore draw the following diagram :




Now suppose we are also told that every C is a B. So we add the additional information by shading the appropriate area, and end up with this diagram :



How should this be interpreted and what should we conclude? Half of the green tick is in a shaded region. What does that mean? Give yourself a minute to think about it before you read on ...

The answer is actually quite simple. The tick indicates that something is both A and C, and it occupies two separate regions. The left hand side region represents things that are A, B and C. The right hand side region represents things that are A and C but not B. Since the tick crosses these two regions, it indicates that there is something either in the class represented by the left region or in the class represented by the right region (or both of course). Shading tells us that there is nothing in the class represented by the right region. So whatever that exists according to the tick must be in the class represented by the left region. In other words, we can conclude that something is A, B, and C. In effect then, shading


“moves” the tick into the left region since it tells us that there is nothing on the right. The above diagram is therefore equivalent to the following one :



So here is a general principle you should remember:


A truncated tick within a region R counts as a complete tick in R if part of the tick is in R and all other parts not in R are in shaded regions.

Exercise #1




Is the Statement “something is either B or C” true according to the diagram?

Exercise #2




Is the diagram consistent with the statement “Everything is B or not C, or both”?

Exercise #3




Is the diagram consistent with the statement “Something is B and not A”?

Exercise #4



Is the diagram consistent with the statement “Everything is A or C”?

Exercise #5



What does the diagram tell us?

VII. Syllogism

We now see how Venn diagrams can be used to evaluate certain arguments. There are many arguments that cannot be analysed using Venn diagrams. So we shall restrict our attention only to arguments with these properties:

- The argument has two premises and a conclusion.
- The argument mentions at most three classes of objects.
- The premises and the conclusion include only statements of the following form: Every X is Y, Some X is Y, No X is Y. Here are two examples :

(Premise #1) Every whale is a mammal.

(Premise #2) Every mammal is warm-blooded.

(Conclusion) Every whale is warm-blooded.

(Premise #1) Some fish is sick.

(Premise #2) No chicken is a fish.

(Conclusion) No chicken is sick.

These arguments are sometimes known as *syllogisms*. What we want to determine is whether they are *valid*. In other words, we want to find out whether the conclusions of these arguments follow logically from the premises. To evaluate validity, we want to check whether the conclusion is true in a diagram where the premises are true. Here is the procedure to follow:


Draw a Venn diagram with 3 circles.

Represent the information in the two premises.

Draw an appropriate outline for the conclusion. Fill in the blank in “If the conclusion is true according to the diagram, the outlined region should.”


See whether the condition that is written down is satisfied. If so, the argument is valid. Otherwise not.

1. Example #1




Let us apply this method to the first argument on this page :


Step 1 : We use the A circle to represent the class of whales, the B circle to represent the class of mammals, and the C circle to represent the class of warm-blooded animals.



Step 2a : We now represent the information in the first premise.
(Every whale is a mammal.)



Step 2b : We now represent the information in the second premise. (Every mammal is warm-blooded.)



Step 3 : We now draw an outline for the area that should be shaded to represent the conclusion. (Every whale is warm-blooded.) This is the red outlined region. We write: “If the conclusion is true according to the diagram, the outlined region should be shaded.”

Step 4 : Since this is indeed the case, this means that whenever the premises are true, the conclusion must also be true. So the argument is valid.

2. Example #2


Let's go through another example:

Every A is B.

Some B is C.

Therefore, some A is C.

We now draw a Venn diagram to represent the two premises:




In the diagram above, we have already drawn a Venn diagram for the three classes and encode the information in the first two premises. To carry out the third step, we need to draw an outline for the conclusion. Do you know where the outline should be drawn?

3. Example #3


Some A is B.

Every B is C.

Therefore, some A is C.



Step 1: Representing the first premise.




Step 2: Representing the second premise.

Step 3: Add an outline for conclusion

VIII. Limitations of Venn diagrams

Although Venn diagrams can help us reason about classes of objects, they also have many limitations. First of all, the diagrams can become too complicated to deal with if we are reasoning about many classes of objects. So far in our tutorials we have considered Venn diagrams with at most three circles. It is possible to add more bounded regions if we are dealing with more than three classes, but then the resulting diagrams will become rather difficult to handle and interpret. It is very easy to make mistakes when we encode information in such diagrams.

The other problem with Venn diagrams is that they have limited *expressive power*. What this means is that there are many pieces of information that cannot be accurately represented. For example, our system of notation allows us to talk about classes of objects, but not particular individual objects. For example, to say that *a* and *b* are cats and *c* and *d* are not, we might have to introduce new symbols, using dots to represent individuals, as in the diagram below:




However, even with this new notation, there are still other pieces of information that cannot be represented, such as:


- Either Felix is a cat or it is a dog.
- If Peter is taller than Mary then Peter is older than Mary.

Perhaps it might be possible to introduce additional new symbols to represent such ideas. But then the system of Venn diagrams will get really complicated and difficult to use. So now that we know the limitations of Venn diagrams, we should be in a better position to know when they are useful and when they are not.

Exercise #1



Is this a suitable Venn diagram for showing the relationships between four sets of objects?



In the second diagram, there are four overlapping rectangles. Which area corresponds to those items which are A, B and D, but not C?

1 This chapter is taken from <http://philosophy.hku.hk/think/> and are in use under the creative commons license.



Critical Thinking by Brian Kim is licensed under a Creative Commons Attribution 4.0 International License, except where otherwise noted.