How to Think Critically

HOW TO THINK CRITICALLY

STEPHANIE GIBBONS AND JUSTINE KINGSBURY

University of Waikato

This book was published by University of Waikato via the Council of Australian University Librarians Open Educational Resources Collective. The online version is available at https://oercollective.caul.edu.au/howtothinkcritically

Disclaimer

Copyright

Recommended citation: Gibbons, S. & Kingsbury, J. (2024). *How to Think Critically*. University of Waikato. https://oercollective.caul.edu.au/howtothinkcritically

CONTENTS

Part III. Chapter 3: Reconstruction

Part IV. Chapter 4: Non-deductive arguments

INTRODUCTION

Welcome to our interactive online textbook *How to Think Critically.* We hope you will find it both useful and enjoyable to work through! In this section, we introduce ourselves and the book, and give you a chance to see how the "interactive" element of the textbook works.

Justine

Justine Kingsbury teaches critical thinking, philosophy of science, philosophy of mind and aesthetics at the University of Waikato. She is from Raglan but now lives in Onewhero, amidst three cats, one dog, two kunekune pigs, nine sheep and some chooks (and one husband).

Stephanie Gibbons teaches critical thinking and logic at the University of Waikato. She has an interest in Ancient Philosophy, particularly metaphysics. She lives in Hamilton and has a smaller collection of animals – just two chickens.

Stephanie

To contact the authors

Stephanie Gibbons
stephanie.gibbons@waikato.ac.nz
Senior Tutor, Philosophy Programme
Te Whare Wānanga o Waikato – The University of Waikato

Justine Kingsbury
Justine.kingsbury@waikato.ac.nz
Senior Lecturer, Philosophy Programme
Te Whare Wānanga o Waikato – The University of Waikato

The textbook has six chapters, each of them accompanied by questions that allow you to test your understanding of key concepts and to practice skills learnt in the chapter.

Table of Contents

Interactive features

This textbook contains interactive features: the text is interspersed with questions designed to test your understanding, and allowing you to practice skills. To make the best use of the textbook, do the questions as you go. Many questions also contain helpful feedback to tell you what you may have missed.

You can try out the interactive feature below, so that you can see how a textbook using this format works.

The most common sort of question is a multiple choice question like this.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=4#h5p-1

A variation on the multiple choice question occurs where there are several correct answers, and you must select all that apply.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=4#h5p-2

ACCESSIBILITY INFORMATION

We believe that education must be available to everyone which means supporting the creation of free, open, and accessible educational resources. We are actively committed to increasing the accessibility and usability of the textbooks we produce.

Accessibility features of the web version of this resource

The web version of this resource has been designed with accessibility in mind by incorporating the following features:

- It has been optimized for people who use screen-reader technology.
    - all content can be navigated using a keyboard
    - links, headings, and tables are formatted to work with screen readers
    - images have alt tags
- Information is not conveyed by colour alone.

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), and various editable files. Choose from the selection of available file types from the 'Download this book' drop-down menu. This option appears below the book cover image on the eBook's landing page.

Third-Party Content

In some cases, our open text includes third-party content. In these cases, it is often not possible to ensure accessibility of this content.

Accessibility Improvements

While we strive to ensure that this resource is as accessible and usable as possible, we might not always get it right. We are always looking for ways to make our resources more accessible. If you have problems accessing this resource, please contact sgibbons@waikato.ac.nz to let us know so we can fix the issue.

PART I
CHAPTER 1: INTRODUCTION

The aim of this book is to help you to improve your reasoning skills. There are three parts to this: a) getting better at detecting when other people are trying to get you to believe things or do things without really giving good reasons why you should, b) getting better at persuading people yourself, and c) getting better at thinking through the reasons you have for believing the things that you believe, and whether or not they are good reasons.

This chapter has the following sections:

WHAT IS CRITICAL THINKING AND WHY IS IT IMPORTANT?

People often try to persuade us of things: the people who write editorials and letters to the editor, politicians, lecturers, advertisers, evangelists, your parents, your children, your friends. And when someone is trying to persuade you of something, you should think: what reasons have I really been given for believing what this person wants me to believe? Are they good reasons? This is what we mean by critical thinking or critical reasoning.

You already do this to some extent, whether or not you know it. Especially if you spend any time online, you are probably constantly being bombarded with information and arguments. You don't believe everything you're told and everything you read. You can't, in fact, because some of the various things you get told conflict with each other. You are already somewhat selective about which of those things you accept – we are going to give you better tools with which to be selective.

One example of a situation in which critical thinking is required is being on a jury.

When you are on a jury, you are faced with a barrage of evidence in support of different conclusions, and you have to weigh it all up and come to an informed decision. The prosecuting lawyer is trying to convince you that the accused committed the crime; the defense attorney is trying to convince you that the prosecutor hasn't proved her case, that at the very least there is doubt about whether the accused committed the crime. And you have to decide what the right conclusion is. Furthermore it matters a lot what you decide – the future of the accused person depends on it. One of our aims (perhaps the main aim) is to make you better at the kind of critical thinking required in this kind of context – better at evaluating what others say when they're trying to convince you of something.

We also hope that learning to recognise the difference between good and bad reasons will make you better at convincing other people of things you want them to believe. Sometimes *you* are the person who is trying to do the persuading. Perhaps you want to persuade people to vote for the political candidate you support or you're the one trying to save their souls; or perhaps you just have a strong urge to get other people to agree with you about the things you care about. Learning to make a good, logical case will go some way towards making you more persuasive. This is important in the jury example as well. After the jury has heard all the evidence, they retire to deliberate – to discuss the evidence and the law and come to a decision about whether or not the defendant

should be found guilty. At that point, as well as being able to evaluate the reasons your fellow-jurors are putting forward, you will want to make a good job of putting forward your own view about whether the defendant is guilty and why. Perhaps most of your fellow-jurors think it is clear that the defendant committed the crime, but you think there are good reasons to doubt this. Being able to state those reasons clearly and straightforwardly may make a real difference to the outcome.

Note that we're not trying to teach you rhetoric – how to use flowery or emotive language which helps you to persuade people of things even when you don't really have a good case. Rather, we're trying to teach you how to give people genuinely good reasons to believe what you want them to believe or to do what you want them to do. The main thing you need to know in order to do this AND in order to evaluate other people's reasoning is what counts as a good reason to believe something and what doesn't.

Here's a bit of rhetoric, just to illustrate what we're NOT trying to teach you to do.

> *Would it be so terrible if there were no slaughterhouses filled with stinking bloody carcasses? Would our lives be poorer if no animals writhed in steel-jaw leg hold traps? Would we fare so badly by wearing natural fabrics instead of furs and leathers? Who could object to a medical science that advanced prevention and healthy lifestyles instead of vivisecting living animals?*

This is all presentation and no content. Phrases and words

such as "stinking bloody carcasses" and "writhed" are highly emotive rather than strictly factual. We are supposed to be swept along by the repetitive rhythm of the rhetorical questions. We're supposed to think that the obvious answers to them are "NO!", "NO!", "NO!" and "NO-ONE!" But are we given any actual reason to think so? None whatsoever.

We will not be teaching you to talk or write like that, although undoubtedly it's a useful skill and one that those who are in the business of persuading others (for example advertisers, and politicians) would benefit from acquiring. Learning to see through rhetorical flourishes, however, is part of what we hope you will get from working through this book. You should think: "Wait a minute! What is this person trying to convince me of? Have I actually been given any reason to believe the thing they are trying to convince me of?"

A further aim, as well as giving you the tools to evaluate other people's reasoning and to produce good arguments yourself in order to persuade other people, is to encourage you to think critically about your own beliefs and the reasons you have (or perhaps don't have!) for holding them. We all have lots of beliefs that we haven't exactly thought through – we may have religious or moral beliefs as a result of our upbringing, for example, and never have thought about whether we actually have good reasons for them.

We will be focusing a lot on how to tell good reasoning from bad reasoning (and suggesting you should apply this both to other people's reasoning and to your own). A secondary focus

will be how to construct chains of reasoning or arguments which are logical, and therefore more likely to convince people. To some extent these two things go together – if you can tell good reasoning from bad, you're more likely to be able to construct good, logical chains of reasoning yourself.

This book will help you to gain skills, or to further develop skills that you already have to some extent , rather than giving you a lot of factual information that you need to memorise. And the way to acquire skills is to practice them a lot. We will be providing you with exercises that you can work through, with feedback: doing these exercises, paying attention to the feedback, and then doing more exercises will help you to become a critical thinker.

ARGUMENTS

In this book, we use the word "argument" in a slightly technical way. In everyday talk, we tend to call any kind of verbal fight an argument. Once I was driving from Wellington to Auckland with my kids, and just out of Levin I had a fight with my seven-year-old about how far it was to Auckland . They insisted that we were nearly there; I insisted that we weren't. A small chunk of the conversation went like this. "We're nearly there." "No, we're not." "Yes, we are." "No, we're not." "Yes, we are." This is not an argument, in the sense of the word which we'll be using here, although it is a verbal fight. Raff and I were making statements or claims that contradicted each other, but not giving any reasons in support of them.

An argument in the sense in which we'll be using the word is a set of statements consisting of a conclusion and some reasons to believe that conclusion.

If either Raff or I gave reasons in support of our claims, we'd have the beginnings of an argument. They thought we were nearly there because they saw a factory that looked like the Tip Top factory, which is on the motorway coming in to Auckland. I thought we weren't for lots of reasons, one of

them being that we hadn't been driving for nearly long enough to be near Auckland.

My argument

> *It takes at least seven hours to drive from to Auckland, and we've only been driving for an hour and a half, so there's still a long way to go.*

Raff's argument

> *I just saw the Tip Top factory.*
> *When you see the Tip Top factory, you're nearly in Auckland.*
> *So we're nearly in Auckland.*

Notice that now we are not just going head to head ("Yes we are!" "No we're not!" etc). Now each of us is giving reasons to believe the thing that we are asserting: each of us is providing an **argument** in support of our claim.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=47#h5p-3

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=47#h5p-4

You can see that not everything people say or write is an argument. A bit of speech or writing is an argument only if the writer or speaker is trying to convince someone of something by providing reasons to think that it is true.

PREMISES, CONCLUSIONS, AND STATEMENTS

An **argument** is a set of statements, some of which are intended to provide reasons to believe one of the others.

The **conclusion** of an argument is the statement that is being argued for, the thing that the arguer is trying to convince her audience of.

The **premises** are the reasons given for believing the conclusion, the evidence that the conclusion is true.

Statements are simply things about which it makes sense to ask "Is it true or false?" "The sky is blue" is a statement. So is "All cats have tails." Not everything we say is like this. If I say "Please pass the salt", that is a request – it doesn't make any sense to ask whether "Please pass the salt" is true or false. If I say "Shut the door!", I'm giving an order, not making a statement. Again, it makes no sense to ask whether "Shut the door!" is true or not. Questions are not statements either.

An interactive H5P element has been

excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=49#h5p-5

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=49#h5p-6

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=49#h5p-7

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=49#h5p-8

When someone makes a statement, they are claiming that something is true.

Sometimes it is appropriate to ask for evidence for a statement, or reasons why you should believe it. If I assert that marijuana should be legalised, you are unlikely to simply believe me, unless you think anything any university lecturer says is automatically true. You'll probably want me to give you some *reasons* in support of my claim.

When someone provides verbal (written or spoken) evidence for a statement, or reasons to believe a statement, they are giving an *argument* for that statement. We call the statements which contain the evidence the *premise*s of the argument. The statement which the evidence is evidence for – the statement that is being argued for – is called the *conclusion*.

An interactive H5P element has been

excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=49#h5p-9

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=49#h5p-10

How can you tell what the conclusion is? Sometimes there are indicator-words, for example, "so" in the first argument above is a conclusion-indicator: "because" in the second argument above is a premise-indicator. (Note however that not every "because" marks a premise in an argument.)

Some conclusion indicators

- therefore
- so
- we can conclude that
- thus

- hence
- consequently

Some premise indicators

- since
- because
- for
- given that

But sometimes an argument is presented without any of these premise or conclusion indicators, and you just have to think: "What is this person trying to convince me of? What is their point?" The thing they're trying to convince you of is their conclusion.

STANDARD FORM

When arguments are given in ordinary language, they often contain material which is inessential to the argument. Much of what people say or write has a rhetorical purpose, and it can be difficult to see what the premises and the conclusion are. To make it easier to assess arguments, we put them into *standard form*.

An argument in standard form puts the premises first, and the conclusion last. The premises and conclusion are labelled, so they can be referred to easily. A line called an *inference bar* is placed between the premises and the conclusion. The inference bar indicates that the conclusion is intended to follow from the premises above it.

Here's an example.

> *There's no need to stop using fossil fuels unless climate change is real. But climate change is fake. So there's no need to stop using fossil fuels.*

In standard form:

> P1) There's no need to stop using fossil fuels unless climate change is real.
>
> P2) Climate change is fake
>
> ---
>
> C) There's no need to stop using fossil fuels.

Note that the 'So' that indicates the conclusion has been left out of the standard form. That's because its role is to indicate where the conclusion is. It is not part of the content of the statement the arguer is trying to get you to believe. Because there are multiple premises, we number them. It does not matter what order you put the premises in, but in standard form they must occur above the conclusion they are intended to support.

One or more interactive elements has been excluded from this version of the text. You can view them online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=51#oembed-1

You can view the transcript for the video 'Standard Form' in the Appendix.

The aim of laying out an argument in standard form is to strip an argument down to its bare bones, so it can be assessed on its merits as an argument alone. This means that the statements occurring in standard form are often rephrased so that they express their content in a straightforward way.

In ordinary language, people sometimes use questions or

commands to communicate statements. When we put the argument into standard form, we rephrase them as statements.

Suppose someone says

> *University education was free for our parents' generation. Shouldn't it be free for us?*

This is an argument: the person is trying to persuade you of something, and they are giving a reason. The reason is "University education was free for our parents' generation". "Shouldn't it free for us?" has the form of a question, but it's being asked rhetorically. The person who gave the argument is not asking a sincere question: they are using the question to make a point. The statement the arguer intends is "University education should be free for us". When we put the argument into standard form, each premise and conclusion is rephrased so that it is expressed as a statement.

> P1) University education was free for our parents' generation.
>
> ---
>
> C) University education should be free for us.

Rhetorical questions are one common way of making statements in ordinary language. Also common is expressing statements as commands. So, someone might say:

> *University education was free for our parents' generation. Make it free for us!*

Once again, a correct standard form rendition of the argument will rephrase the command as a statement.

> P1) University education was free for our parents'

generation.

C) University education should be free for us.

Later, we will need to assess the truth of statements in arguments. As each statement will have its truth assessed individually, each statement should be complete in itself. Consider the following argument:

Barack Obama was a great president. He was always dignified and composed.

The sentence "He was always dignified and composed" is incomplete, because we don't know, when considering the sentence by itself, who is referred to by "he". To make the statement complete, all we need to do is replace the "he" with the person it stands for. This is clear from the context. So, in standard form:

P1) Barack Obama was always dignified and composed.

C) Barack Obama was a great president.

An interactive H5P element has been excluded from this version of the text. You can view it online here:

https://oercollective.caul.edu.au/howtothinkcritically/?p=51#h5p-11

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=51#h5p-12

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=51#h5p-13

Often, to put an argument into standard form you will need to omit a lot of what is said. Consider this example:

> *Vegans are a bunch of puppy-kissing tree-huggers. They seem to think everybody should be vegan just because they are. Everyone is entitled to their own opinion. Besides, they*

give lousy reasons for being vegan: just as many animals are hurt through harvesting of plants as die in animal farming.

We need to distill the argument from this. Approach such pieces of writing by asking yourself

- what is the person trying to persuade me of?
- what reasons do they give to support that?

Here, the conclusion is not explicitly stated, and you need to work it out. The main reason is clear, however. We can put the argument into standard form as follows.

P1) Just as many animals are hurt through harvesting of plants as die in animal farming.

C) It is not the case that everyone should be vegan.

The rhetoric has been removed, and the argument has been expressed in simple statements.

You may think something is missing from the above argument. You are right. In chapter 3 we will look at how to make such standard form reconstructions more complete. Before we do so, we need to look at one important aspect of what makes a good argument: validity.

CHAPTER 1 QUESTIONS

This section contains questions. Use them to check your understanding!

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=53#h5p-14

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=53#h5p-15

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=53#h5p-16

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=53#h5p-17

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=53#h5p-18

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=53#h5p-19

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=53#h5p-20

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=53#h5p-21

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=53#h5p-22

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=53#h5p-23

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=53#h5p-24

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=53#h5p-25

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=53#h5p-26

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=53#h5p-27

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=53#h5p-28

PART II

CHAPTER 2: DEDUCTIVE ARGUMENTS

You'll remember from Chapter 1 that an argument is an attempt to persuade someone of some conclusion by giving reasons in support of that conclusion.

Deductive arguments are ones which are intended to guarantee the truth of their conclusion.

For example, this is a deductive argument:

P1. All politicians are dishonest.
P2. Winston is a politician.

C. Winston is dishonest.

You can see that if the premises were true, the conclusion would be guaranteed.

Not all arguments try to absolutely conclusively prove their conclusion, and an argument can still be a good argument if it makes its conclusion *very likely* to be true, rather than *certain* to be true.

For example, consider this argument.

P1. 92% of politicians are dishonest.
P2. Winston is a politician.

C. Winston is dishonest.

The premises, even if true, don't *guarantee* the conclusion. But if the premises are true then the conclusion is *very likely* to be true.

The first argument above (with the "All" statement as its first premise) is a deductive argument – it's clear that the arguer intends to conclusively prove the conclusion. The second one, with the percentage claim as its first premise, is not: there's no way you could *guarantee*, with premises like these, that the conclusion was true. (It's always possible that Winston is in the honest 8%.) But the second one is still quite a good argument – it provides quite a good reason (supposing the premises are true) to think that Winston is dishonest.

Sections:

INTRODUCTION TO ARGUMENT EVALUATION

There are two things you always need to think about in deciding whether or not an argument is a good argument:

1) Are the premises and the conclusion connected in such a way that the premises, if they were true, would provide good reasons to accept the conclusion?

Here is an example of an argument in which the premises are not connected to the conclusion in this way:

P1. All mothers are female.

P2. Justine is female

C. Justine is a mother

If you think about the premises, you can probably see that even if they are true, they still don't give you any reason to believe that Justine is a mother. It would be different if P1 said "All females are mothers" – from that and P2, you certainly could conclude that Justine is a mother (though then there would be a new problem with the argument: "All females are mothers" is clearly not true). But if you're given the argument as stated, you should think "Hang on a minute. Just because

Justine is female and all mothers are female, that doesn't mean that Justine is a mother. She might be, or she might not be. These premises aren't enough reason to think she is."

2) Are the premises actually true?

Here is an example of an argument in which the premises are properly connected to the conclusion but the premises aren't true:

P1. All rugby players sing opera.
P2. Kiri Te Kanawa is a rugby player.

C. Kiri Te Kanawa sings opera.

If it were true that all rugby players sing opera and that Kiri Te Kanawa is a rugby player, then it would just have to be true as well that Kiri Te Kanawa sings opera. The connection between the premises and the conclusion is the strongest possible connection – the premises, if they were true, would guarantee the truth of the conclusion. So this argument passes the first test for being a good argument.

But it doesn't pass the second test. It's an argument in which, if the premises were true, the conclusion would be guaranteed. However, the premises aren't true. All rugby players sing opera? Surely not! Kiri is a rugby player? Don't think so. If one or more of the premises is false, we shouldn't accept the conclusion on the basis of those premises.

In evaluating an argument, we always have to consider these two things. We have to think about what the connection is between the premises and the conclusion. To do this, we have

to forget about whether the premises are actually true and think about whether, if they *were* true, they would give us good reason to believe the conclusion.

And we have to decide whether we think the premises are true. Obviously, if we don't believe in the premises then, even if the conclusion follows logically from the premises, we don't have to accept the conclusion. On the other hand, if we think that the premises are true, and if the conclusion follows logically from the premises, then in the interests of consistency we should believe in the conclusion as well.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=57#h5p-29

VALIDITY

For the rest of this chapter we are going to be talking specifically about evaluating *deductive* arguments – non-deductive arguments will come later, in Chapter 4.

A deductive argument is one which is intended to guarantee the truth of its conclusion. The terms we use in evaluating deductive arguments are **validity/invalidity** and **soundness/unsoundness**.

First, validity. A valid argument is one in which *if* its premises were all true, its conclusion would have to be true as well. It doesn't matter (for validity) whether in fact the premises are true. All that matters for validity is that there should be a connection between the premises and the conclusion such that if the premises *were* true, the conclusion would also have to be true. A valid argument is an argument in which it is impossible for the premises to be true and the conclusion false.

The validity of an argument is independent of whether the premises are in fact true. Thus you don't need to know anything about the subject that the argument is about in order to judge whether or not it is valid. To say that an argument is valid is to say something about its *structure*, not to say anything about its content. When we talk about validity, we are talking

about the first of the two argument evaluation tasks above: we are talking about what the connection is between the premises and the conclusion.

A **valid** deductive argument is one in which *if* all the premises were true, the conclusion would also have to be true.

For example

P1. No men are mothers.
P2. Some students are men.

C. Some students are not mothers.

and

P1. All rugby players sing opera.
P2. Kiri Te Kanawa is a rugby player.

C. Kiri Te Kanawa sings opera.

Remember that when what you are considering is an argument's validity, it doesn't matter whether the premises are actually true. So it doesn't matter, for the moment, whether it's true that no men are mothers or that all rugby players sing opera. What matters is what connection (if any) there is between the premises and the conclusion. A valid argument has the strongest possible connection between premises and conclusion – so strong that if the premises were all true, the truth of the conclusion would be guaranteed.

So in the first example above, to see why the argument is valid, think: suppose it's true that no men are mothers and that

some students are men. Then, must it also be true (on that supposition) that some students are not mothers?

The answer is that supposing those premises to be true, it must also be true that some students are not mothers. So the argument is valid.

Doing the same for the Kiri Te Kanawa example: Suppose that it was true that all rugby players sang opera and that Kiri was a rugby player. Then the conclusion would have to be true as well: it would have to be true that Kiri sings opera. The argument is valid.

You may be thinking at this point, "But that's stupid! We all know that it's not true that all rugby players sing opera! So how can the argument be valid?"

Bear in mind that validity is not the *only* thing you have to take into consideration in deciding whether an argument is a good argument or not: it also matters whether the premises are true. The Kiri argument is valid, but it is still not a good argument. We will get onto this issue in a bit.

In everyday language, the word 'valid' is often used to mean 'true' or 'reasonable'. In philosophy generally, and in this course, 'valid' has a technical meaning. An argument which is valid is one where it is impossible for the premises to all be true and the conclusion false.

An interactive H5P element has been

excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=59#h5p-30

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=59#h5p-31

Here are a number of different ways of saying what a valid argument is. They all amount to the same thing – you can use whichever one or ones help you to understand validity.

A **valid** argument is one in which:

- it is impossible to have all the premises true and the conclusion false at the same time.
- the conclusion logically follows from the premises.
- if the premises were all true, the truth of the conclusion would be guaranteed.
- if the premises were all true, the conclusion would also

have to be true.

When you ask whether or not a particular argument is valid, you are talking about its structure, not about its content. Consider the following argument.

P1. All adlers are bobkins.
P2. All bobkins are crockers.

C. All adlers are crockers.

You can tell it's valid even if you don't know what adlers, bobkins or crockers are. *I* don't know what they are, so I have no idea whether it's *true* that all adlers are bobkins or that all bobkins or crockers. Nevertheless, I know the argument is valid: because of its structure, if the premises *were* true, the conclusion would have to be true too.

You cannot usually tell, from the truth or falsity of the premises and conclusion of an argument, whether it is valid or invalid. A valid argument can have false premises and a false conclusion, or false premises and a true conclusion, or true premises and a true conclusion. An invalid argument can too. There's just one kind of case in which you *can* tell about an argument's validity or invalidity from the truth or falsity of its premises and conclusion, and that is when you have an argument which has true premises and a false conclusion. A deductive argument with true premises and a false conclusion has to be invalid, because a valid argument by definition can never have true premises and a false conclusion.

So when you are deciding whether or not an argument is valid, don't think about whether the premises and conclusion are true. Instead, *imagine* or *suppose* that the premises are true, and then think about whether that would mean that the conclusion also had to be true.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=59#h5p-32

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=59#h5p-33

An interactive H5P element has been excluded from this version of the text. You

can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=59#h5p-34

You might be wondering at this point why we should care about validity. Since validity has nothing to do with whether or not the premises are true, what's the point of it? We shouldn't accept the conclusion of an argument on the basis of its premises if its premises are obviously false, as they are in the Kiri Te Kanawa argument or in the "sisters and brothers" argument above. So why bother pointing out that the argument is valid, since it's obviously a really bad argument?

In those cases, perhaps in real life contexts you wouldn't need to. There are two conditions a deductive argument needs to satisfy in order to be any good: it has to be valid, and it has to have true premises. Once we notice that this one has false premises, we can already tell it's a bad argument, whether or not it's valid.

But it is important to be able to tell whether an argument is valid in other cases. One kind of case in which it matters is when a deductive argument has all true premises. That's not enough to make it a good argument: you need to check whether or not it's valid as well. For example, suppose someone argued like this:

P1. February is the next month after January.

P2. Grass is green

_____

C. Snow is white.

Although the premises are true, they don't connect properly to the conclusion – in fact they don't connect to the conclusion at all. So they provide no reason whatsoever why you should believe the conclusion. This shows that just having true premises is not enough (not *nearly* enough!) to make an argument a good argument.

Another kind of case in which being able to assess validity is of practical importance is when others disagree with you about the falsity of the premises. In some contexts it is useful to be able to point out that, even though you think the premises are false, that *even if they were true* the argument would be no good.

COMMON ARGUMENT PATTERNS

Some argument types are so common they have their own names. Learning to recognise these patterns will help you recognise valid arguments.

This section introduces four common argument patterns, and some simple variations on them.

All of the argument patterns in this section are valid.

- Modus ponens
- Modus tollens
- Disjunctive syllogism
- Hypothetical syllogism
- Some notes on Conditionals and generalisations

Modus ponens

Consider the following argument:

P1) If Rover is a dog, then Rover is a mammal.

P2) Rover is a dog.

C) Rover is a mammal.

It has the following pattern:

If p then q

p

Therefore q

The letters 'p', 'q', 'r' etc. are a traditional way of representing statements. Any statement can be inserted in the place of 'p' and 'q', and the resulting argument will be valid. Learning to recognise some common argument forms makes it much easier to identify valid (and invalid) arguments.

Here's another argument with the same pattern.

P1) If Winston Peters is a Cabinet Minister, then New Zealand First must have entered into a coalition agreement with either the National Party or the Labour Party.

P2) Winston Peters in a Cabinet Minister.

C) New Zealand First must have entered into a coalition agreement with either the National Party or the Labour Party.

This argument has longer statements, but the basic pattern is the same. The pattern is identified as Modus ponens, so it must be valid.

An interactive H5P element has been excluded from this version of the text. You can view it online here:

https://oercollective.caul.edu.au/howtothinkcritically/?p=61#h5p-35

One or more interactive elements has been excluded from this version of the text. You can view them online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=61#oembed-1

You can view the transcript for the video 'Modus ponens' in the Appendix.

Modus tollens

Modus tollens is another very common valid argument form. It has the following pattern:

If p then q
Not q
Therefore not p

The first premise claims that if p occurs, then q must also occur. The second premise points out that q hasn't occurred. So it has to follow that p hasn't occurred either. Why? Because

if p had occurred, then q would also have occurred. And we know q hasn't occurred.

It may be easier to see with an example.

P1) If Hillary Clinton has won the last election, then she would be president.

P2) Hillary Clinton is not president.

C) Hillary Clinton did not win the last election.

It's clear that it is not possible for this conclusion to be false while these premises are true. This is a valid argument.

Try this one. Remember, you're looking to see whether the pattern of Modus tollens applies.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=61#h5p-36

Try another one:

An interactive H5P element has been

excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=61#h5p-37

Disjunctive syllogism

A 'disjunction' is a complex statement where two statements are joined with an 'or' (or another word serving the same role). A **disjunctive syllogism** is a valid argument with the following form:

p or q
Not p
Therefore q

This argument form is valid because the initial premise dictates that one of the two options must hold, and the second premise asserts that one does not hold. It follows that the other must hold.

Here is an example.

P1) Chess is the most challenging board game or Monopoly is the most challenging board game.
P2) Chess is not the most challenging board game.

C) Monopoly is the most challenging board game.

It does not matter whether it is what is before the 'or' or what is after the 'or' which is denied. But it must be denied.

You can practice applying the pattern in the following questions.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=61#h5p-38

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=61#h5p-39

A disjunctive syllogism is often expressed using an 'either... or... construction. For instance:

P1) Either there will be a recession, or house prices will continue to rise.

P2) House prices will not continue to rise.

C) There will be a recession.

Sometimes a disjunctive syllogism uses 'either' along with 'or', and sometimes it doesn't. It doesn't change the force of 'or' either way. 'Either' is generally used rhetorically, to emphasise the contrast between the two options. There is more on 'either' in the next section on Structural fallacies.

Hypothetical syllogism

A hypothetical syllogism creates a 'chain' of conditional claims. So long as the links of the chain occur in the right way, where each leads to the next, the intermediate links can be omitted. Here's an example:

P1) If housing prices continue to rise, then rents will continue to rise.

P2) If rents continue to rise, then rental accommodation will become unaffordable for the working poor.

C) If housing prices continue to rise, then rental accommodation will become unaffordable for the working poor.

It has the general form:

If p then q
If q then r
Therefore if p then r

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=61#h5p-40

When checking an argument form, the order of the premises is irrelevant. That is because validity treats the premises as a collection of claims. You are welcome to change the order of the premises if it makes it easier for you.

One or more interactive elements has been excluded from this version of the text. You can view them online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=61#oembed-2

You can view the transcript for the video 'Order of premises' in the Appendix.

Conditionals and generalisations

Conditionals

Several of the basic argument patterns above use **conditional** claims.

A conditional is an 'if... then...' statement. The 'if...' part of the statement is called the '**antecedent**', and then 'then...' part is called the '**consequent**'.

Conditionals are sometimes expressed in a different order. The 'antecedent' is the 'if...' clause no matter what order the parts are presented in.

So, the English sentence 'If Borka is a goose, then Borka is a bird' means the same thing as 'Borka is a bird if she's a goose'.

The same conditional can be expressed in more ways than this. One version of a conditional that people find especially tricky is 'only if'. Suppose there is a sign in the university carpark which says: "Staff permit holders only". This means "You can only park here if you are a staff permit holder". Think about what this involves. It does not mean that if you are a staff permit holder you must park there. It doesn't forbid staff permit holders from parking in other places. What it means if that if you are not a staff permit holder, you must not park there. That is equivalent to "If you park here, then you are (must be) a staff permit holder".

What this shows is that the order of antecedent and consequent in a conditional statement is very important, and

they cannot simply be reversed. To see this, let's return to the goose example.

It is true that "if Borka is a goose, then Borka is a bird". However, it is not the case that "If Borka is a bird then she is a goose". Some birds are not geese. So, "Borka is a bird only if she is a goose" is false. The 'only if' claim which is equivalent to "If Borka is a goose then she is a bird" is "Borka is a goose only if she's is a bird".

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=61#h5p-41

Conditional statements can also be expressed using 'unless'. "If Borka is a goose then she's a bird" is equivalent to "Borka is not a goose unless she's a bird". We often use 'unless' in contrast to a 'not' claim. So, someone might say "I won't babysit for you unless you pay me", which means the same as "If I babysit for you then you are (must be) paying me".

An interactive H5P element has been

excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=61#h5p-42

Conditionals and generalisations

There is also an important relationship between conditionals and generalisations. The reason why 'if Borka is a goose, then Borka is a bird' is true, is because the generalisation 'All geese are birds' is true.

Any hard generalisation can be expressed as a conditional. 'All geese are birds' can be expressed as 'if something is a goose, then it is a bird'. This means that the basic argument patterns which use conditionals all have forms which use generalisations instead.

All As are Bs
x is an A
Therefore x is a B

is a variation on Modus ponens, using a generalisation in the place of a conditional.

Here is an example of a hypothetical syllogism, using generalisations instead of conditionals:

P1) All squares are quadrilaterals.

P2) All quadrilaterals are polygons.

_____

C) All squares are polygons.

Any argument with this form will be valid.

Here are some for you to try:

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=61#h5p-43

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=61#h5p-44

An interactive H5P element has been

excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=61#h5p-45

BASIC STRUCTURAL FALLACIES

A fallacy is a bad argument. These fallacies are similar to the valid argument patterns above, and are frequently confused with them. These argument forms, however, are invalid. It is a good idea to practise recognising the difference: you will need to pay close attention to the patterns.

This section looks at three common structural fallacies:

- The fallacy of affirming the consequent
- The fallacy of denying the antecedent
- Disjunctive fallacies

The fallacy of affirming the consequent

The fallacy of affirming the consequent is an invalid argument which is often mistaken for Modus ponens or Modus tollens. It has the following form.

If p then q
q
Therefore p

In this fallacy, the consequent of the first premise is affirmed

in the second premise. Such an argument isn't valid. The first premise claims that when p occurs, q must also occur. But it does not claim that the occurrence of q guarantees the occurrence of p.

The problem is easy to see with a couple of examples.

P1) If Bernie Sanders is the president of the US, then Hillary Clinton is not president of the US.

P2) Hillary Clinton is not the president of the US.

C) Bernie Sanders is the president of the US.

The first premise is true. It is true that 'If Bernie Sanders is the president of the US then Hillary Clinton is not the president of the US'. There can only be one president of the US at a time, and if it is Sanders, then it isn't Clinton. The second premise is also true. Hillary Clinton is not the president of the US. But it certainly doesn't follow from that that Bernie Sanders is the president.

Here's another example.

P1) If this shape is a square then its sides are equal in length.

P2) This shape's sides are equal in length.

C) This shape is a square.

It is true that any square will have sides equal in length. But it is perfectly possible to have a shape with equal sides which is not a square. (An equilateral triangle, for instance.) So this conclusion does not follow from the premises.

It is relatively easy to see the problem with these examples. Sometimes examples which affirm the consequent are harder to spot. It helps to look at the form of the argument carefully. It often helps to work out the form using letters instead of statements. This prevents you from being distracted by what you know (or believe) is true.

Remember, the fallacy of affirming the consequent is an *invalid* argument form.

Can you correctly identify the argument form?

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=63#h5p-46

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=63#h5p-47

The fallacy of denying the antecedent

The fallacy of denying the antecedent is another invalid argument which is often mistaken for Modus ponens or Modus tollens. It has the following form.

If p then q

Not p

Therefore not q

Here an example of the argument: you should be able to see that this one is invalid.

P1) If it's wrong to eat meat then it's wrong to eat human beings.

P2) It's not wrong to eat meat.

C) It's not wrong to eat human beings.

It must be true that if it's wrong to eat meat then it's wrong to eat humans. But it does not follow from it being permissible to eat meat that it's permissible to eat people. There can be compelling reasons to not eat people even if it turns out that eating other sorts of meat is okay.

When assessing these it's important not to get distracted by what you know or believe. Consider this example:

P1) If Hillary Clinton is president then the president is a woman.

P2) Hillary Clinton is not president.

C) The president is not a woman.

The conclusion does not follow from the premises. It is possible to imagine a world in which it was true that Hillary Clinton was not president but *where some other woman was president*. That is a world where the premises are true and the conclusion is false. Therefore the argument is not valid.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=63#h5p-48

Disjunctive fallacies

A disjunctive syllogism has the following form. It is valid.

p or q
Not p
Therefore q

This form, however, is not valid:

p or q
p
Therefore not q

There are, however, situations in which we would be willing to accept arguments of the second form. Here's an example.

P1) Donald Trump or Joe Biden will win the next presidential election.
P2) Joe Biden will win the next presidential election.

C) Donald Trump will not win the next presidential election.

You may well look at this argument and say that it is valid: if the premises are both true, then the conclusion has to be true also.

The reason for this is that you know there can only be one winner of the next election. It is not possible for both Trump and Biden to win. So, if Biden wins, Trump cannot. Here, the real meaning of P1 is 'Donald Trump or Joe Biden will win the next election, *but not both'*.

Sometimes in English we use the word 'or' to mean 'or, and both are possible', and sometimes we use or to mean 'or, but not both'. The first of these is called an 'inclusive or', and the second is called an 'exclusive or'.

The form given for a disjunctive syllogism above is valid on either interpretation of 'or'. But this form

p or q
p
Therefore not q

is only valid if the 'or' is exclusive. That is, it can only be valid if it is not possible to have both p and q. This is something which can only be determined by considering the

meaning of the claim made. It cannot be determined by looking at the form of the argument.

When you try to assess this sort of argument, think to yourself 'Is it *possible* for both p and q to occur?'. If it is possible, then the argument isn't valid. If it's impossible, then the argument is valid. However, it may help to add 'but not both' to the disjunctive premise.

Sometimes people claim that an 'either... or...' construction is used to show that both options cannot hold: that is, an 'either... or...' construction indicates an exclusive 'or'. It might be useful if we used 'either... or...' in this way in English, but we don't. If I said to you 'Bring either beer or wine to the party; I don't mind', and you turned up with both, you would be rightfully put out if I then said 'You brought both. You can't come in. I said to bring either one or the other.' The presence of the word 'either' does not tell you whether an exclusive 'or' or an inclusive 'or' is being used. You must use your common sense for that. A good rule of thumb is if you're not sure, suppose the 'or' is inclusive.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=63#h5p-49

SOUNDNESS

A deductive argument is **sound** if (and only if) it meets two conditions.

1. It is valid.
2. The premises are true.

A valid argument with true premises is guaranteed to have a true conclusion. This means sound arguments are very useful.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=65#h5p-50

An interactive H5P element has been excluded from this version of the text. You can view it online here:

https://oercollective.caul.edu.au/howtothinkcritically/?p=65#h5p-51

We will return to sound arguments in detail later, when we learn more about assessing truth in chapter 5. For the moment you should simply be aware that a sound argument is a valid argument with true premises.

CHAPTER 2 QUESTIONS

These questions allow you to test your understanding of the chapter 2 skills.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=67#h5p-52

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=67#h5p-53

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=67#h5p-54

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=67#h5p-55

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=67#h5p-56

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=67#h5p-57

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=67#h5p-58

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=67#h5p-59

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=67#h5p-60

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=67#h5p-61

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=67#h5p-62

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=67#h5p-63

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=67#h5p-64

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=67#h5p-65

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=67#h5p-66

PART III

CHAPTER 3: RECONSTRUCTION

When we encounter arguments in real life, they are usually not neatly set out in standard form for us to evaluate. Sometimes they're embedded in lots of extraneous material, including rhetoric; often not all of the premises are explicitly stated; sometimes the premises or the conclusion are unclear. This happens because often an arguer has other aims as well as the aim of giving a good argument – for example, a blog post or an opinion piece in the newspaper might aim to grab our attention with a catchy headline or opening sentence and then keep us interested and engaged (and perhaps outraged) all the way to the end, and a clear and precise argument in standard form may not do that job.

Because of this, when we want to evaluate a real-life argument, often we first have to reconstruct it. The full process of reconstructing an argument involves more than putting it into standard form. We "clean up" the argument so that it is as clear and straightforward as possible, so that we are ready to assess it.

Reconstruction treads a line between interpreting the arguer in a charitable way, and doing too much work for the

arguer. We want to clarify what the arguer is doing, and present their argument in the optimum way. But it still needs to be their argument which is presented, and not some other (better) argument. We look for what the arguer likely intended, but we don't add new or better reasons for what they have said.

The result of an argument reconstruction is an argument which is ready to be assessed for soundness.

Chapter 3 Sections

IMPLICIT PREMISES

When people give arguments in ordinary language, they often leave parts of their arguments out. Often this is because something is so obvious it can be safely assumed that others will accept it, and so it doesn't need to be explicitly stated. Consider the following argument.

> *My pet Squeaky is a mouse, and all rodents have teeth that never stop growing. So, Squeaky's teeth will never stop growing.*

There is an unstated assumption here. That is that mice are rodents. Without assuming this, the conclusion of the argument would not validly follow from the premises.

The difficulty with leaving premises unstated is that sometimes the unstated premise is not obvious or easily accepted, but is in fact a highly controversial claim. For this reason, we make any implicit premises explicit when reconstructing arguments. This means that when we assess the argument we can properly assess each premise as true or false.

With the above example, we begin by putting the argument into standard form.

P1) My pet Squeaky is a mouse.

P2) All rodents have teeth that never stop growing.

C) Squeaky's teeth will never stop growing.

We then note that the argument is invalid. We could make it valid, however, by adding a premise, like so:

P1) My pet Squeaky is a mouse.

P2) All rodents have teeth that never stop growing.

P3) All mice are rodents.

C) Squeaky's teeth will never stop growing.

The argument is now valid. And, this is a sensible addition to the argument: it's clearly something that the arguer intended, even though it wasn't explicitly said.

Sometimes an implicit premise is left out by the arguer because it is so obvious it is hardly worth saying. However, sometimes an unstated premise is doing a lot of work in the argument, and that isn't evident because it hasn't been explicitly stated. Sometimes the unstated premise is obviously false, or highly controversial. By exposing implicit premises, and making them explicit, we're better positioned to assess the argument.

Consider this argument:

Co-sleeping is risky for the baby. So no one should do it.

An initial reconstruction might look like this:

P1) Co-sleeping with a baby carries a risk of harm to the baby.

C) No one should co-sleep with a baby.

What is the missing premise here? What is needed to make the argument valid? To make the argument valid, a connection needs to be made between the risk of harming the baby, and what shouldn't be done. So to make the argument valid, we could add an implicit premise such as this:

P1) Co-sleeping with a baby carries a risk of harm to the baby.

P2) No one should do anything with a baby that carries a risk to the baby.

C) No one should co-sleep with a baby.

This is the minimum that is required to make the argument valid. The arguer must have something like this in mind, otherwise the conclusion of the argument wouldn't follow. Here, the connecting premise is doing a lot of work in the argument, and it is false. It cannot be true that no one should do anything that would put a baby at risk. If that was true, people would never to be able to take a baby in a car, or an aeroplane, or do very much at all with them. Living a life free of risk would be paralysing, for a baby or for anyone else.

It's likely that the arguer really meant that the risk of co-sleeping is an unacceptable risk. However, given that their argument doesn't make an attempt to evaluate risk, or to explain what degree of risk would be acceptable, to adjust their argument in this way would be to do too much work for them. In the absence of any attempt to explicitly link the premise to

the conclusion, there is not much we can do but provide the minimally necessary connection, and assess it.

Working out what premise needs to be added to an argument to make it valid is tricky. You need to think about how validity works, and how to connect together what has been provided to ensure that the conclusion follows. The following video gives you some hints to get you started. It's a good idea to watch it before attempting the questions.

One or more interactive elements has been excluded from this version of the text. You can view them online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=71#oembed-1

You can view the transcript for the video 'Connecting Premises' in the Appendix.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=71#h5p-67

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=71#h5p-68

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=71#h5p-69

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=71#h5p-70

Normative conclusions

Arguments with normative conclusions deserve special mention. They are very common, and they frequently have implicit premises which need to be made explicit.

A **normative** claim is a statement says what *should* or *ought* to happen. In contrast, a **descriptive** claim is a statement which says how things *are*. So, the statement "Mount Everest is the highest mountain in the world" is a descriptive statement. It describes a current feature of the world. "Climbers should seek permission before climbing Mount Everest" is a normative statement, as is "Fewer people should climb Mount Everest".

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=71#h5p-71

An interactive H5P element has been excluded from this version of the text. You

can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=71#h5p-72

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=71#h5p-73

Conclusions are often normative, because arguments are often trying to persuade people about how things ought to be, or about what ought to happen.

In order to be valid, an argument with a normative conclusion must have at least one normative premise. No valid argument can have only premises which describe the way the world is, and conclude something about how things should be.

Consider this argument:

P1) Some people are finishing their schooling unable to read.

C) We should implement a more comprehensive literacy programme in our schools.

The conclusion is a normative one. The factual claim that some people are unable to read can never be enough to support the normative claim. No amount of information about how things are can ever, on its own, support the claim that things should be different. So we can see, merely by noting that the conclusion is normative, that a normative premise is needed for validity.

Here the argument can be rendered valid by adding a conditional:

P1) Some people are finishing their schooling unable to read.

P2) If some people are finishing their schooling unable to read, then we should implement a more comprehensive literacy programme in our schools.

C) We should implement a more comprehensive literacy programme in our schools.

Inferring a normative conclusion from descriptive premises is such a common type of argument failure that it has its own name: it is called "the fallacy of deriving ought from is". However, it is generally easily avoided by adding the necessary normative premise.

Using conditionals and generalisations as connecting premises

Any argument, no matter how far the premises appear to be removed from the conclusion, can be made valid. Sometimes, when there is no obvious link between the premises and the conclusion, the easiest way to make the argument valid is by adding a conditional claim. The definition of a valid argument is "if the premises are true, then the conclusion must be true also". So, by constructing a conditional that says "if [other premises], then [conclusion]", and adding that conditional as a further premise, any argument can be rendered valid. We can call such a conditional a "corresponding conditional" (because it is the conditional that corresponds to the argument).

Here's how it works.

Suppose someone gives an argument that says

> *Fish can't ride bicycles. Therefore the moon isn't made of green cheese.*

In standard form:

P1) Fish can't ride bicycles.

C) The moon is not made of green cheese.

It might look, at first glance, as if there is no way to connect these two ideas. They are completely unrelated. But, we assume that the arguer intended there to be a connection, because we assume the arguer meant to give a good argument.

To connect them, we form a corresponding conditional: we say "if the premise, then the conclusion". Here, that gives us

P1) Fish can't ride bicycles.

P2) If fish can't ride bicycles, then the moon is not made of green cheese.

C) The moon is not made of green cheese.

This is a valid argument.

Whether the argument is a good argument or not depends on whether or not P2 is true. Using a corresponding conditional to make an argument valid shifts the burden of assessing the argument to a particular premise.

A corresponding conditional will make any argument valid, but it is not always the best way of making an argument valid. You will still (ultimately) need to assess the truth of the conditional, and that is not always very easy to do. Sometimes a good technique is to start with a corresponding conditional, and then to think about whether there is some generalisation which makes that conditional true. So, take the argument

P1) Squeaky is a mouse.

C) Squeaky is a mammal.

The corresponding conditional which makes this argument valid is "If Squeaky is a mouse, then Squeaky is a mammal". But what makes *that* claim true is the generalisation "All mice are mammals". So, we can use the generalisation to connect P1 to the conclusion:

P1) Squeaky is a mouse.

P2) All mice are mammals.

C) Squeaky is a mammal.

Whether it is better to use a generalisation or a conditional will depend partly on how specific the connection needs to be. Consider this argument:

P1) It is Friday night.

C) Jerry will end up drunk tonight.

What sort of connection could be used here? The corresponding conditional is "If it is Friday night, then Jerry will end up drunk". There are a wide range of generalisations that cover this conditional. Here are some options:

1. Everyone gets drunk every night.
2. Everyone gets drunk on Friday nights.
3. Jerry gets drunk every night.
4. Jerry gets drunk every Friday night.

The first and second of these would make the argument valid. However, they're not charitable connections to choose, because they are clearly false. We don't want to make the argument worse than it needs to be. While the third option is possible, the fourth seems a better option. Given that we know nothing but what is claimed in the premise, and we need to connect it to the conclusion, it seems likely that the

arguer wants to appeal to something such as this. It is more general than the corresponding conditional. But there would be no reason for the arguer to present P1 as a reason for the conclusion unless a generalisation such as the fourth option were true. So this generates a good reconstruction of the argument as valid:

P1) It is Friday night.

P2) Jerry gets drunk every Friday night.

C) Jerry will end up drunk tonight.

If you're unsure how to make an argument valid, it can help to begin with a corresponding conditional, and then to make that claim more general. The limit on how general you should make it is the limit of what is reasonable or believable. If there is no reasonable generalisation available, it may be better to leave it as a corresponding conditional.

One or more interactive elements has been excluded from this version of the text. You can view them online here: https://oercollective.caul.edu.au/howtothinkcritically/?p=71#oembed-2

You can view the transcript for the video 'Finding a more general connecting premise' in the Appendix.

Although a corresponding conditional can be used to make any argument valid, it should not always be used. Sometimes there is a more obvious option for a connecting premise. For example, consider this argument:

P1. All mice have tails.

C. Minnie has a tail.

The corresponding conditional is "If all mice have tails, then Minnie has a tail." But a simpler premise to add as P2, which like the corresponding conditional would do the job of making the argument valid, would be "Minnie is a mouse."

CLARIFYING ARGUMENTS

In Chapter 1, we noted that when putting arguments into standard form, we need to make sure each statement is clear and complete. There are further points to keep in mind about keeping an argument clear and easy to assess.

Phrasing

In ordinary language, people refer to the same thing in a number of different ways in a short space. This prevents them from sounding repetitive. When reconstructing we ensure the alternative phrasing "matches up" so that the same phrase is used throughout the argument. This makes it easier to spot validity and invalidity. Consider the following argument.

> *Factory farming of pigs should be banned! The conditions the pigs live in are disgusting. No farm animal should be kept in such deplorable conditions. Prohibit the practice!*

A first go at putting this argument into standard form might go like this:

P1. The conditions the pigs live in are disgusting.

P2. No farm animal should be kept in such deplorable

conditions.

C. Factory farming of pigs should be banned.

The conditions are described as "disgusting" in P1, and as "deplorable" in P2. (The conclusion was also presented in two ways, but as the same claim is being made in both the first and last sentence, we only need to keep one of them.) Although slightly different phrasing has been used in different places, it's clear that this is the same point is being made. We can also note that the context makes it clear that P1 is about factory-farmed pigs, and not pigs in general. We adjust the phrasing to make it "match up" as follows:

P1. The conditions in which factory-farmed pigs are kept is unacceptable.

P2. No farm animal should be allowed to be kept in unacceptable conditions.

C. The factory-farming of pigs should not be allowed.

This is valid. It's now clearer that it's valid, because the repetition of particular phrases makes it easier to see that the same things are being talked about.

Quantifiers

One of the ways that people frequently leave statements incomplete is by leaving out quantifiers. Quantifiers are words which indicate *how many* or *how much*. Quantifiers are often

dropped from claims, and this can cause problems when assessing arguments. So, when someone says "Students drink too much", do they mean *all* students, *most* students, or *some* students? When we reconstruct we work out what the arguer most likely intended, and make the quantifier explicit.

When quantifiers are left out, it sometimes generates confusion. Consider this argument:

> *Men need to look to their behaviour! Men at the Oxfam charity evening were groping waitresses and making lewd comments. That is totally unacceptable.*

Such arguments are sometimes responded to with the cry "Not all men!". Such a response can be avoided by making the quantifier explicit. We can give a charitable rendition of the argument as follows:

> P1. **At least some men** at the Oxfam charity evening were touching waitresses uninvited and speaking to them inappropriately.
>
> P2. It is not acceptable to touch someone uninvited or to speak to them inappropriately.
>
> P3. Anyone who is engaged in unacceptable behaviour should change their behaviour.
>
> ---
>
> C. **At least some men** should change their behaviour.

This argument is valid – the "at least some men" from P1 is carried through to the conclusion. And it is not susceptible to the "Not all men!" criticism, because it is not making a claim about all men.

Where possible, make your quantifiers explicit.

IRRELEVANT MATERIAL

Often arguments are embedded in other material which is not strictly speaking part of the argument. Sometimes this is an oversight on the part of the arguer – they haven't noticed that some of their claims are irrelevant to their conclusion. But sometimes it's not – perhaps the arguer has put the extra material in to provide the context for the the argument, or to catch the attention of the reader, or to fulfil any of the many possible purposes the arguer might have which are different from the purpose (which they must also have, since we've said they are an *arguer*) of providing reasons in support of some conclusion.

One thing that you may have to do in reconstructing an argument is get rid of extraneous or irrelevant material – claims which are not premises in the argument because they are not relevant to the conclusion. Now we are going to talk a bit about what it means to say that a claim is *irrelevant* to a conclusion.

Suppose a student comes to my office and makes this argument (this, with minor variations, happens quite often):

> *I deserve an A for this course. If I don't have an A average, I won't get into second-year law, and if I don't get an A for this course I won't have an A average.*

Their conclusion is that they deserve an A for the course. But if you think about it, none of the rest of what they say backs up the claim that they deserve an A. The kind of information that would back up the claim that they *deserve* an A would be to do with the quality of their work. For example, if they bring me their assignments and show me that actually their answers are much better than the grader has realised, so they should get As rather than Bs for their assignments, that would be relevant to the conclusion that they deserve an A for the course. But what they actually say has no bearing on the question of whether they deserve an A. Rather, they are giving reasons why they *need* an A. But "I need an A for this course" is a quite different conclusion. (The student may be being sneaky. That you need an A for the course is not a reason for me to give you one – grades are determined by quality of work, not by need. So overtly trying to get me to give you an A by telling me how much you need one is unlikely to work. On the other hand, if you can get your lecturer to think that you *deserve* an A, they might give you one. The problem is that if the lecturer has any critical thinking skills, they will be able to see that you're not actually giving any reason to think that you *deserve* an A.)

What it means to say that a statement is **irrelevant** to a particular conclusion is that whether that statement is true or false makes no difference whatsoever to whether we should believe the conclusion; it has no bearing on the conclusion at all.

When you are deciding about relevance and irrelevance, you first have to identify the conclusion. A statement is not relevant or irrelevant on its own – it's only relevant or irrelevant to some particular conclusion.

Note that being irrelevant to a conclusion is not the same thing as being false. You can have irrelevant statements that are nevertheless true – for example, in our example, everything the student says about second-year law, etc, might be true, but that wouldn't stop it from being irrelevant. And you can have false premises that are nevertheless relevant to the conclusion.

Here is an example. My colleague Clare works part-time, and she has a sign on her office door that says: "Clare is not here in the mornings." For the purposes of these questions, you should assume this bit of background information is true: My colleague Clare works part-time, and she is at work every afternoon but never in the mornings.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=75#h5p-74

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=75#h5p-75

Now it is 3 pm, and the sign says "Clare is out of her office having afternoon tea", but in fact Clare has returned from her break and is in her office with the door closed marking essays. I want to know whether or not she is in her office.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=75#h5p-76

An interactive H5P element has been excluded from this version of the text. You

can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=75#h5p-77

The Clare example shows (we hope) the difference between relevance and truth.

Now that we've talked about what it means for a claim to be *irrelevant* to a conclusion, we will look at an example of a passage of text in which some of the claims are relevant to the conclusion and some aren't. Part of the job of reconstructing such an argument is to identify which claims are relevant, and put them into a standard form version of the argument as premises, while leaving out the claims that are irrelevant to the conclusion.

Here's the example. My grandmother used to say to me, when I didn't eat all the food on my plate: "Eat your dinner! You're being wasteful. There are children starving in Africa!" I was always puzzled about what the starving children in Africa had to do with me eating or not eating my dinner – after all, it's not as though the food I was wasting was food that could otherwise have been eaten by them. I think I was right to be sceptical. The argument really goes something like this (P2 isn't stated, but I think it's implied):

P1. Not eating all the food on your plate is wasteful.

P2. We shouldn't be wasteful.

C. You should eat all the food on your plate.

I have not included "There are children starving in Africa" as a premise, because I don't think it's relevant to the conclusion – I have cleaned up the argument by leaving it out.

MORE COMPLEX ARGUMENTS

Extended arguments

The arguments we have looked at so far have been short and simple. It is very common, however, for arguments to have several steps. Here's an example:

> *If the minimum wage is raised, then shops will need to put their prices up in order to pay their workers. If shops put their prices up, then we will all pay more for goods, and the general cost of living will rise. So, if the minimum wage goes up, the general cost of living will rise. We should not increase the cost of living for those who can least afford it. Therefore the minimum wage should not be raised.*

The final conclusion of this argument is "The minimum wage should not be raised". There is another conclusion in the argument, however. In the third sentence a conclusion is drawn. That conclusion is then used as a reason for the final conclusion. So it is operating as both a conclusion, and as a premise. We can call such conclusions, which are also premises, "intermediate conclusions". To see how that works, here's the argument in standard form:

P1. If the minimum wage is raised, then shops will

need to put their prices up in order to pay their workers.
P2. If shops put their prices up, then we will all pay more for goods, and the general cost of living will rise (for everyone).

C1/P3. If the minimum wage goes up, the general cost of living will rise (for everyone).
P4. We should not increase the cost of living for those who can least afford it.

C. The minimum wage should not be raised.

An argument can have any number of intermediate conclusions. An argument has only one final conclusion: that is the ultimate thing the arguer wishes to persuade you of. We still number premises sequentially, but in an extended argument the intermediate conclusions are also numbered.

When assessing the validity of an extended argument each inference is assessed. (An inference occurs whenever the arguer proposes that a conclusion follows from premises. In standard form, an inference occurs every time an inference bar is used.) If even one inference is invalid, the argument as a whole is invalid.

Consider this argument:

> *If a foetus is a person then abortion is murder. To be a person, something must be self-aware. But a foetus isn't self-aware. It follows that a foetus is not a person. And so abortion is not murder.*

In standard form we get an extended argument as follows:

P1. To be a person, something must be self-aware.
P2. A foetus is not self-aware.

C1/P3. A foetus is not a person.
P4. If a foetus is a person then abortion is murder.

C. Abortion is not murder.

The first inference, from P1 and P2 to C1, is a valid inference. But the second inference, from C1 and P4 to C2, commits the fallacy of denying the antecedent. It is not a valid inference. So this argument as a whole is invalid.

To identify an intermediate conclusion it is best to start at the end of the argument and work backwards. So, begin by asking what the arguer is (ultimately) trying to persuade you of. This is the final conclusion. Then, ask what reasons are given for that. Once you have those reasons, ask yourself if those reasons are supported by any further reasons. If a reason is supported by a further reason, it is operating as an intermediate conclusion. It can often help to sketch out the structure of the argument. For examples of how to do this, see the section on argument trees below.

An interactive H5P element has been excluded from this version of the text. You

can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=77#h5p-78

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=77#h5p-79

Dealing with longer texts

Even relatively short arguments can have intermediate conclusions. Most of the arguments you will encounter in everyday life are not this short. Most arguments occur within longer texts which include background information, irrelevant material, pieces of rhetoric, anecdotes and stories, repetition, and (sometimes) rambling. If you try to go through a long piece of text sentence by sentence, it is easy to end up overwhelmed, and to get bogged down. Instead, try the following:

- Read the text all the way through, and then put it out of sight.
- Ask yourself "what was the main point the person was trying to persuade me of?". That is the final conclusion.
- Ask yourself "what were the main reasons the person gave for the conclusion?". These are the main premises.
- Ask yourself whether anything was said to support those main reasons.

You're now putting together the main line of argument. It is likely that none of what you have noted down follows the wording of the original. That is fine. It is not the precise words the arguer uses that matter, but the ideas which they communicate.

Once you have have noted down the main line of argument in this way, you can return to look at the original argument to see if you missed anything important. But it is still important not to get bogged down by particular phrasing, or by the detail of the argument. It is more important to present the basic point to others in your reconstruction. You can't easily do this if you have 37 premises. Instead, make a fair and charitable reconstruction of the main points, and then assess those.

ARGUMENT TREES

Argument trees, also known as argument diagrams or tree diagrams, are a useful way of mapping out the structure of arguments. Because argument trees are inherently visual, you will find the videos a more useful way of learning about how these work.

One or more interactive elements has been excluded from this version of the text. You can view them online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=79#oembed-1

The transcript of the video 'Argument diagrams 1' can be viewed in the Appendix.

One or more interactive elements has been

excluded from this version of the text. You can view them online here: https://oercollective.caul.edu.au/howtothinkcritically/?p=79#oembed-2

The transcript of the video 'Argument diagrams 2' can be viewed in the Appendix.

CHAPTER 3 QUESTIONS

Test your chapter 3 skills!

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=81#h5p-80

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=81#h5p-81

An interactive H5P element has been

excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=81#h5p-82

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=81#h5p-83

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=81#h5p-84

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=81#h5p-85

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=81#h5p-86

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=81#h5p-87

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=81#h5p-88

PART IV

CHAPTER 4: NON-DEDUCTIVE ARGUMENTS

So far we have been discussing deductive arguments: arguments that are intended to conclusively prove their conclusions. However there are many arguments that give us good reasons to believe their conclusions even though they are not deductively valid. In this chapter we consider these non-deductive arguments.

INTRODUCTION TO NON-DEDUCTIVE ARGUMENTS

The famous fictional detective Sherlock Holmes is always talking about *deducing* his conclusions from the evidence, but actually the kind of reasoning he engages in is not deductive reasoning.

Here is an extract from the Sherlock Holmes story "A Scandal in Bohemia."

> Dr Watson visits Holmes after a long absence. Holmes figures out that Watson has started practising medicine again, and that he has been out in bad weather lately, and that he has an incompetent servant, even though Watson hasn't told him any of these things.
>
> *"My eyes tell me that on the inside of your left shoe, just where the firelight strikes it, the leather is scored by six almost parallel cuts. Obviously they have been caused by someone who has very carelessly scraped round the edges of the sole in order to remove crusted mud from it. Hence, you see, my double deduction that you had been out in vile weather, and that you had a particularly malignant boot-slitting specimen of the London slavery. As to your practice, if a gentleman walks into my rooms smelling of iodoform, with a black mark of nitrate of silver upon his right forefinger, and*

> *a bulge in the side of his top hat to show where he has secreted his stethoscope, I must be dull indeed if I do not pronounce him to be an active member of the medical profession*."

In this passage, Holmes draws three conclusions:

1. Watson has been out in bad weather.
2. Watson has an incompetent servant.
3. Watson has started practicing medicine again.

The evidence presented for 1 and 2 is that the leather on the inside of Watson's left shoe is scored by six almost parallel cuts. The evidence presented for 3 is that Holmes smells of iodoform, has a black mark of nitrate of silver upon his right forefinger, and has a bulge in the side of his top hat.

Even if these are good reasons to believe the conclusions, surely they do not guarantee them. There are other logically possible reasons why Watson might have six cuts on his shoe, smell of iodoform, etc. There is an unstated premise that the hypothesis that Watson has started practising medicine again is the best explanation for the iodoform smell, black mark on his finger and bulge in the side of his top hat. And that this is the best explanation of the observed facts is intended as a reason to believe it. Such arguments can never guarantee the truth of their conclusion – it's always possible that some other explanation is actually the correct one. So they are not deductive arguments. Nevertheless, such non-deductive arguments can provide very good reasons to believe their

conclusions. We will be discussing a number of different kinds of non-deductive arguments in this chapter, concluding with arguments like the Sherlock Holmes one, which is called an inference to the best explanation.

One way to signal that your argument is non-deductive is to put [Probably] in front of the conclusion – this shows that you are not intending to guarantee that the conclusion is true, only to make it likely (probable). And one test of whether an argument is deductive or non-deductive is to think about whether it would make sense to put a [Probably] in front of the conclusion. For example, if someone argued that since all mice have tails and Minnie is a mouse, Minnie has a tail, it would make no sense to put [Probably] in front of the conclusion "Minnie has a tail." The premises, if true, make it *certain* that Minnie has a tail, not just probable. On the other hand, if the argument instead went "Almost all mice have tails, and Minnie is a mouse, so Minnie has a tail," it makes very good sense to put a [Probably] in front of "Minnie has a tail," and that's a sign that the argument is non-deductive.

Non-deductive arguments are never valid, and therefore never sound. But, since they're not trying to be, we don't tend to complain about them being invalid and unsound– that's the wrong standard to hold a non-deductive argument to. The terms we use when evaluating a non-deductive argument are "strong" and "cogent".

Strength plays the same role in the evaluation of non-deductive arguments as **validity** plays in the evaluation of

deductive arguments. To say that a non-deductive argument is **strong** is to say that its premises, if they were true, would provide good reason to believe the conclusion, although (since it's not a deductive argument) they do not *guarantee* the truth of the conclusion. Like validity, strength is to do with how strong the connection is between the premises and the conclusion. Like validity, strength has nothing to do with whether or not the premises are actually true.

An argument is **cogent** if it is strong **and** all of its premises are true.

We ask the same two questions when evaluating non-deductive arguments as we did when evaluating deductive arguments:

1. What is the connection between the premises and the conclusion?
2. Are the premises true?

In the case of non-deductive arguments, we answer the first question by talking about **strength**. When we talk about **cogency**, we are answering both questions.

There is an important difference between validity and strength. Validity is all-or-nothing. A deductive argument is either invalid or invalid; it cannot be a little bit valid or partly valid or almost valid. Strength is not like this: it is a matter of degree. Some non-deductive arguments provide almost complete support for the conclusion, some provide *a bit* of

support for the conclusion, and some provide hardly any or no support for the conclusion.

Consider the following three arguments:

Argument One:

P1. 96% of politicians are dishonest.

P2. Winston is a politician.

C. Winston is dishonest.

Argument Two:

P1. 75% of politicians are dishonest.

P2. Winston is a politician.

C. Winston is dishonest.

Argument Three:

P1. Most politicians are dishonest.

P2. Winston is a politician.

C. Winston is dishonest.

The first is very strong, the second is less strong, the third is even less strong. Strength is a matter of degree.

Summary

A **strong** argument is one in which the premises provide a lot of support (although not conclusive support) for the conclusion: if the premises were all true, the conclusion would be likely to be true.

Strength is a matter of degree (unlike validity).

A **cogent** argument is a strong argument which does in fact have all true premises.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=85#h5p-89

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=85#h5p-90

Here's an everyday example of non-deductive reasoning. There used to be a film critic on National radio called Jonathan Dennis. He was a rather harsh critic – he didn't like many movies, and in particular, he very rarely liked a commercial, Hollywood movie. So, since I do like some such movies, when Jonathan said that a particular movie was bad, I didn't necessarily abandon my plans to go and see it. But if Jonathan

raved about a movie that I wasn't planning to go to, I was quite likely to decide to go and see it after all.

My reasoning went like this:

P1. In cases W, X, Y and Z, when Jonathan liked a movie, I liked it too.

P2. There haven't been any cases when Jonathan liked a movie and when I saw it I didn't like it.

P3. Jonathan liked movie P.

[Probably] C. I will like movie P as well.

Do you think I was reasoning well, or badly?

Some of you might have said that four cases (movies W, X, Y and Z) are not enough to generalise from. Certainly, on the basis of this amount of evidence, I wouldn't go so far as to say that any movie that Jonathan Dennis liked, I would also like. If I was thinking of making any such strong claim, I'd want to test it further. If I was going to be scientific about it – if I was going to set about trying to test the claim – I'd go to more movies recommended by Jonathan Dennis and see if I liked them.

But think about the context. All I'm trying to decide about is whether to go to a movie. The most I stand to lose if I get it wrong is one evening and $15. I think the evidence on which I am deciding to go to the movies is good enough, under the circumstances. That is to say, even if the argument doesn't make the probability of the conclusion *very* high, it's good enough for my purposes.

If something more serious was at stake, we would want better evidence, more cases, something which gave us reason to think the conclusion had a *very* high likelihood of being true. Suppose my child is diagnosed with cancer and I am choosing between conventional medical treatments such as chemotherapy and some alternative treatment like "quantum vibrational therapy," which involves passing an electric current through the body. I should want a lot more than four cases in which the alternative therapy has been successful before I'll think I have good reason to use it. In a case in which so much is at stake, I need much stronger reasons to believe the conclusion than in a case in which I am merely deciding whether or not to go to a movie.

PROBABILISTIC ARGUMENTS

In this section we will cover:

- Argument patterns
- Soft generalisations
- Extended probabilistic arguments
- A note on "Probably"

But first, some general comments on probabilistic arguments.

Probabilistic arguments occur where the likelihood of the conclusion can be clearly established given the premises. They are the type of non-deductive argument which most closely resembles the way deductive arguments work.

Consider this argument:

P1. All sheep in New Zealand live on farms.

P2. Alice is a sheep in New Zealand.

C. Alice lives on a farm.

Suppose for a moment that Alice is a New Zealand sheep (that is, suppose that P2 is true). The argument is valid. However, it cannot be sound. The first premise is a "hard" generalisation – it does not allow of any exceptions. As a hard

generalisation about all sheep in New Zealand, P1 is false. There are undoubtedly some rogue sheep. There will be a few who have escaped into the bush, and there are probably a few sheep kept as pets who don't live on farms. So although it is valid, this argument is unsound.

We could change P1 to a "soft" generalisation that has a better chance of being true. A soft generalisation makes a general claim about a group, but allows that there are some exceptions. So if P1 used the statement

> Nearly all sheep in New Zealand live on farms

then it would be true.

But the argument would no longer be valid:

> P1. Nearly all sheep in New Zealand live on farms.
>
> P2. Alice is a sheep in New Zealand.
>
> ———————
>
> C. Alice lives on a farm.

In this argument the premises do not guarantee the conclusion. It is *possible* for the premises to be true but the conclusion false, because it is possible that Alice is one of the few rogue bush-sheep, or a pet.

This sort of argument isn't valid, but it can be very useful. The premises fall short of guaranteeing its conclusion, and so it isn't valid. But it does provide strong support for its conclusion. The truth of the premises are sufficient to show that the conclusion is *probably* true.

This sort of argument isn't a failed deductive argument: it isn't intending its conclusion to follow with certainty. We

can mark this in the argument frame, by including the word "Probably" before the conclusion, like so:

P1. Nearly all sheep in New Zealand live on farms.

P2. Alice is a sheep in New Zealand.

[Probably] C. Alice lives on a farm.

Arguments of this sort will be stronger or weaker, depending on how probable the premises make the conclusion. Some such arguments have premises that make their conclusions very probable, and so are very strong.

P1. There are 99 black marbles in this bag and one white marble.

P2. In my fist is a marble randomly selected from the bag.

[Probably] C. The marble in my fist is black.

Here we can see that it is 99% probable that the marble in my fist is black. That makes this a very strong argument.

It's important to note that the statement "The marble in my fist is black" is still either true or false. It cannot be 99% true. It is either 100% true or 100% false. The 99% applies to the probability that it is true, and not to truth itself.

We can change the probabilities in such arguments by changing the proportions of marbles:

P1. There are 75 black marbles in this bag and 25 white marbles.

P2. In my fist is a marble randomly selected from the

bag.

[Probably] C. The marble in my fist is black.

This conclusion is still probable. This non-deductive argument is weaker than the previous one, but it is still strong enough to be a useful argument.

With the marbles example it is easy to accurately measure the degree of probability of the conclusion. Most ordinary probabilistic arguments lack this degree of precision.

P1. Most university students do not have children.

P2. Betty is a university student.

[Probably] C. Betty does not have children.

Here the conclusion is probable, but we can't assign a precise degree of probability to the conclusion.

Argument patterns

The same sorts of argument patterns can occur in probabilistic non-deductive arguments as occur in deductive arguments.

The argument "No mammals lay eggs. Perry is a mammal. So Perry does not lay eggs" is a valid argument. It follows the general pattern of Modus tollens. (If you can't see why, try converting the generalisation expressed in the first premise as a conditional. Now the first premise reads "If something is a mammal then it doesn't lay eggs.") But the first premise of this argument is false. There are three species of mammal that lay

eggs: the most well-known is the platypus. So, we can soften the generalisation in the first premise:

P1. Hardly any mammals lay eggs.

P2. Perry is a mammal.

[Probably] C. Perry doesn't lay eggs.

This follows the Modus tollens pattern, except that it uses a soft generalisation instead of a hard generalisation in the first premise. It is a non-deductively strong argument.

It's important to remember that a fallacious argument pattern cannot be improved by weakening the generalisation being made. Here's an example to make the point clearer.

Consider this argument:

P1. All geese are birds.

P2. Borka is a bird.

C. Borka is a goose.

The basic pattern of this argument is the fallacy of affirming the consequent (using a generalisation instead of a conditional). There are plenty of birds which are not geese, and Borka could be one of those.

This argument cannot be improved by weakening the generalisation in P1. That would give us an argument like this:

P1. Most geese are birds.

P2. Borka is a bird.

[Probably] C. Borka is a goose.

This is a weak argument. Borka is not likely to be a goose in virtue of being a bird. Once again, Borka could be some other type of bird. So this argument commits a non-deductive version of the fallacy of affirming the consequent. That basic problem with the structure of the argument cannot be avoided by one of the premises being a soft generalisation instead of a hard one.

This might seem obvious, but you do see people give versions of this type of argument. Consider this one, which appeared in the media not that long ago.

P1. Nearly all terrorists are Muslim.

P2. The person sitting next to me on the plane is Muslim.

[Probably] C. The person sitting next to me on the plane is a terrorist.

When such arguments are criticised, the people giving them sometimes respond by saying "I wasn't saying all terrorists are Muslim, only most of them", or by saying "I'm not saying I'm certain that they're a terrorist, just that it's likely". But these things are not the problem with the argument. Even if it were true that nearly all terrorists are Muslim (which it isn't), the conclusion would not be likely to be true. That's because this argument commits the fallacy of affirming the consequent. It cannot make it likely that the person sitting next to me is a terrorist. The argument patterns we learnt in the previous

chapter remain relevant for probabilistic arguments using soft generalisations.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=87#h5p-91

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=87#h5p-92

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=87#h5p-93

Kinds of soft generalisation

Any statement which makes a claim about a group or category of things can be said to be a generalisation. **Hard generalisations** are things like "All", "None," "Always" and "Never". **Soft generalisations** include "Almost all," "Almost none," "Many," "Most," and "Some". Some soft generalisations are useful in probabilistic arguments and some are not.

Remember that the aim is to show that the conclusion is probable: that is, to have any strength at all, the argument has to show that the conclusion is more likely to be true than not.

Consider this argument:

P1. The majority of men beat their wives.

P2. Aristotle is a man.

[Probably] C. Aristotle beats his wife.

This argument has some strength, although not much. If the premises were true, then the conclusion would be more likely to be true than not. However, P1 is clearly false. So the argument cannot be cogent.

We might try to improve the argument by softening the generalisation in P1 further, in an attempt to make it more likely to be true. An attempt to do so might give us something like this:

P1. Many men beat their wives.

P2. Aristotle is a man.

[Probably] C. Aristotle beats his wife.

P1 is now more likely to be true. But this argument is weak. The premises don't provide good reason for accepting the conclusion. The "many" doesn't provide a proportion of all men which could make the conclusion probable. Words such as "many" don't tell us proportions. It just means that there are at least several men who do so. It's important to think about whether the generalisation is sufficiently robust to make the conclusion probable.

Extended probabilistic arguments

Just as deductive arguments can occur as extended arguments, so can non-deductive ones. We need to watch how any probability which occurs in the argument effects the probability of the final conclusion.

With an extended argument with a single soft generalisation, the probability of the conclusion will reflect the degree of probability in the soft generalisation:

P1. Nearly all university students write assignments on computers.

P2. Betty is a university student.

[Probably] C1. Betty writes her university assignments on a computer.

P3. Everyone who writes assignments on a computer

can read.

[Probably] C. Betty can read.

This is a strong argument. The probability of C is the same as that of C1, and the degree of probability of C1 comes from the soft generalisation in P1.

But each soft generalisation in an extended argument will further dilute the probability of the final conclusion.

Consider this argument:

P1. Most university students hand in their assignments.

P2. Conrad is a university student.

[Probably] C1. Conrad hands in his assignments

P3. Most students who hand in their assignments pass their courses.

[Probably] C. Conrad passes his courses.

Here, the inference from P1 and P2 to C1 is not particularly strong. It is further weakened by the soft generalisation at P3. By the time the final conclusion is reached, the probability given to the conclusion by the premises is low. This argument is not a strong argument.

If you can't see why the dilution is causing a problem from that example, consider this one, where the problem is more obvious:

P1. Most of those currently in the university library

are university students.

P2. Conrad is currently in the university library.

[Probably] C1. Conrad is a university student.

P3. Most university students drink in the evenings.

P4. It is evening.

[Probably] C. Conrad is drinking.

It's unlikely (although not impossible) that Conrad is drinking in the library. Even if all the premises of this argument were true, the conclusion is not likely. That's because the group of people who are likely to be in the university library in the evening are likely to be different people than those who are drinking.

Sometimes the generalisations in an extended argument will be strong enough to make the final conclusion still probable, and sometimes they will not. There is no precise way to determine the probability of the conclusion when imprecise quantifiers such as "nearly all" and "few" are used. Instead, watch the number and type of generalisations made, and make a judgement call about whether the probability of the conclusion has been diluted too much.

A note on the use of "Probably"

When putting non-deductive arguments into standard form, we often insert "[Probably]" before the conclusion to indicate

that the argument is intended to be non-deductive. It's put in square brackets to indicate that it is not part of the conclusion itself, and it is not part of the argument itself. It merely indicates the type of argument being used. This is often helpful, but it's important to note that it does not indicate anything about the success of the argument. Nor can inserting "Probably" improve a poor argument. Consider this argument:

> Nearly all dogs have four legs. Fido is a dog. So Fido has four legs.

This is a strong argument. It is strong regardless of whether "Probably" is placed before the conclusion or not. So

> P1) Nearly all dogs have four legs.
>
> P2) Fido is a dog.
>
> _____
>
> C) Fido has four legs

is a strong argument.

Further, inserting "Probably" before the conclusion does not show that an argument is strong, and will not improve a poor argument. Consider this argument:

> P1) Nearly all dogs have four legs.
>
> P2) Fido has four legs.
>
> _____
>
> [Probably] C) Fido is a dog.

This is not a strong argument. (It affirms the consequent.) The presence of "Probably" cannot change this. You should think of the "Probably" as a useful way to indicate that an

argument is non-deductive. But it does not tell you anything about the success of the argument.

ENUMERATIVE INFERENCES

Imagine a turkey living happily on a turkey farm. Every morning the farmer brings corn for it to eat, and it doesn't take much more than that to make a turkey happy. One morning the farmer approaches and the turkey (let's suppose) thinks happily "Hooray, breakfast." Supposing the turkey is reasoning at all, it's reasoning non-deductively: On every morning up until now the farmer has brought corn, so today he'll be bringing corn again. But alas, it's Christmas morning, and the turkey is making a terrible mistake when he runs happily towards the farmer to get fed, because this time the farmer is bringing an axe instead.

The turkey was reasoning, if it was reasoning at all, pretty well: it was reasoning from true premises, and ones which gave it quite a lot of reason to believe its conclusion. Nevertheless, its conclusion was false. Moral: no matter how good your non-deductive argument, it's still possible for your conclusion to be false.

Someone who is reasoning like this is taking a (large) number of cases of which they have experience, and inferring that a pattern which has occurred will continue to occur. They have, if you like, collected some data, and they are

extrapolating from that data to formulate a conclusion. Inferences of this type are sometimes called "inductive inferences". But because this is not the only type of induction, we have chosen to call these "**enumerative inferences**". That is because a number of cases are collected, and, on the basis of that list of cases, a conclusion is reached about a new case.

Enumerative inferences are different from probabilistic arguments. Consider this probabilistic argument from above:

P1. There are 75 black marbles in this bag and 25 white marbles.

P2. In my fist is a marble randomly selected from the bag.

[Probably] C. The marble in my fist is black.

In this argument, the proportions of the contents of the bag are known, and, because this is a mathematical example, the degree of probability of the conclusion can be precisely calculated. (It is 75% likely that the marble in my fist is black.)

Now suppose I have a bag of marbles, and I know that there are 100 marbles in it. I know nothing about what colour the marbles are, however. I draw out the first 99 marbles, and they are all black. On this basis, I conclude that the 100th marble will also be black. My argument looks like this:

P1. Marble 1 is black.

P2. Marble 2 is black.

P3. Marble 3 is black.

⋮

P99. Marble 99 is black.

[Probably] C. Marble 100 is black.

I can't assign a precise degree of probability to that conclusion. There are infinite possibilities for what (shade of) colour the remaining marble could be. But it certainly seems more reasonable for me to suppose that it is black, given the contents of the bag so far, than to suppose it is some other colour.

In ordinary life, we reason like this all the time. When I suppose that the sun will come up tomorrow morning, I'm extrapolating from many cases of the sun coming up. This has happened every day of my life so far, and I expect it to continue. I suppose that if I get hit by a bus, I will get hurt, and I suppose this on the basis of what usually happens when people get hit by buses, and on the basis of what has happened to me in the past when I have been hit by large heavy objects. You might think that all I need to work out what would happen if I was hit by a bus is to apply the laws of physics, but even the belief that the laws of physics will continue to apply is justified through an enumerative inference. Such arguments are very important, and very useful.

Not all enumerative inferences are strong. They are often difficult to assess. Consider once again the marbles case above. When I know that there are 100 marbles in the bag, and the first 99 have been black, it seems reasonable to conclude that the 100th marble will be black. But what if I didn't know how

many marbles were in the bag? What if I have only drawn 10 marbles? Can I still (justifiably) conclude that the next marble will be black?

There are a number of things to take in consideration when assessing an enumerative inference.

- How big is the sample?

The more data that is collected, the stronger the enumerative inference. That is why an inference about the colour of the next marble is stronger when you have tested 99 marbles than when you have tested only 9. The sample of times the sun has come up in the morning is huge, and we feel very confident that there will be a tomorrow.

- How big is the sample compared to the total population?

If I know that there are a million marbles in the bag, and I've tested 99, I'll feel less confident about the next marble I pull out than if there are 100 marbles in the bag and I've tested 99.

The size of the total population can also vary depending on what the conclusion is claiming. Sometimes a conclusion is about the next case alone. So, consider:

Argument 1

P1. The sun has come up every day of my life.

[Probably] C. The sun will come up tomorrow.

I feel very confident about this conclusion. The sample size is all the days of my life, and the total population with which the argument is concerned is all-the-days-of-my-life +1. The sample is a very big proportion of the total population.

Compare:

Argument 2

P1. The sun has come up every day of my life.

[Probably] C. The sun will come up every day for the rest of my life.

and

Argument 3

P1. The sun has come up every day of my life.

[Probably] C. The sun will come up every day forever.

In argument 2, the total population is unknown, but I'm going to be optimistic and suppose I'm about halfway through my life. That means I'm extrapolating from known cases to about the same number of cases again.

In argument 3, the conclusion is so broad that it's unlikely to be true. We know that the world will end some day, so there will, one day, be a day that's the last day. A conclusion which goes too far beyond its sample ends up having a sample which is a very small proportion of the total population.

- How was the sample collected?

Suppose I am given 10 bags of marbles, with 10 marbles in them each. If I take 10 marbles from the first bag, and they are all black, I have *some* reason for thinking all the marbles in all the bags are black, but it's not a particularly good reason. If I take 1 marble from each bag, and they are all black, then I have a much better reason for thinking all the marbles are black.

Generally speaking, the more data you have collected, the stronger the enumerative inference. The more random the data collection, the stronger the enumerative inference. If you find yourself rejecting almost all enumerative inferences, your standard for reasonableness is probably too high. We do use these inferences all the time. And, in fact, it would be impossible to function without them. You'd have no reason not to step in front of a bus.

Here are some for you to try.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=89#h5p-94

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=89#h5p-95

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=89#h5p-96

It's important to note that the possibility of being wrong is not sufficient grounds for rejecting an enumerative inference. This possibility of inferring a false conclusion from true premises is a feature of all non-deductive arguments. Consider again the turkey from the beginning of this section. The turkey is justified in his conclusion even though one day he will be wrong. One day the sun won't come up. That doesn't mean you are unjustified in believing it will come up tomorrow.

ARGUMENTS FROM SAMPLES

Whenever we sample or survey some (but not all) members of a group and then draw a conclusion about the group as a whole, we are reasoning non-deductively. We've had an example of this already:

> *A nationwide poll of a random sample of thousands of home-owners revealed that 70% of them are opposed to increases in social welfare payments. Therefore, roughly 70% of the adult population of New Zealand opposes such increases.*

In that case, remember, there was an obvious problem with the argument: all of people surveyed were home-owners, but the conclusion drawn is about the adult population, not just about home-owners. The sample was not *representative*.

Suppose we do a better survey: instead of only asking home-owners, we draw our sample randomly from the adult population of New Zealand. And suppose the results come out like this: 55% of the sample of thousands of adult New Zealanders oppose increases in social welfare payments. We conclude that 55% of all adult New Zealanders oppose such increases.

This is a stronger argument than before: now the sample

from which we are generalising is a sample that is representative (so far as we can tell) of the wider population we are generalising to. Note, though, that the argument is still non-deductive. Unless you poll every single member of the wider population (in which case you are no longer arguing from a sample) the conclusion that what is true of the sample is also true of the wider population is not guaranteed.

The other thing we need to take into consideration, as well as the representativeness of the sample, is the *size* of the sample. In the example above, if we had only surveyed 10 randomly-selected New Zealand adults, rather than thousands, then we should definitely not generalise from the results of the survey to a conclusion about adult New Zealanders in general – the sample size is much too small.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=91#h5p-97

ANALOGY

An analogy is when you point out the similarities between two different things.

If I say "the mind is like a computer: it takes certain inputs, processes them, and then spits out results", I'm drawing an analogy – I'm saying that these two different things – the mind and a computer – are actually alike in some way. But I'm not giving an argument – I'm not drawing any conclusion from my analogy. I'm just drawing parallels, maybe trying to make you think about the mind in a different way, illustrating something about how the mind works.

An **argument by analogy** is when you point out the similarities between two or more things and then draw some conclusion.

Suppose there is a particular car that I'm deciding whether or not to buy. It's a ten-year-old Honda Civic with 75,000 kms on the clock, it has a little rust but not much, and it's only had one owner and she drove it carefully and got it serviced regularly.

And suppose I reason like this:

> My last car was a Honda Civic, and when I bought it it was 10 years old had 75,000 kms on the clock, it has a little rust but not much, it had had only one previous

owner, and that owner had looked after it well. And my last car served me well – I drove it for 5 years and it gave me no trouble and needed very little in the way of repairs. So this car, being so similar to my last, will probably likewise serve me well – I should buy it.

P1. Car A was a Honda Civic, ten years old, 75,000 kms, only a little rust, well looked after, and it served me well for 5 years each and required very little in the way of repairs.
P2. Car B is a Honda Civic, ten years old, 75,000 kms, only a little rust, well looked after.

C. Car B will last 5 years and hardly need any repairs.

General structure of arguments by analogy:

A has characteristics W, X, Y and Z
B has characteristics W, X, and Y

B will have characteristic Z as well.

Things to consider in evaluating an argument by analogy:

- How similar are the things being compared?
- Are the similarities relevant? (e.g. if the similarities mentioned were all things like colour, which is surely

irrelevant to how well the car goes, it would not be a good argument.)

- Are there any relevant differences between the things being compared?
- How many similar cases are we dealing with? (e.g. If I had had 3 cars which were similar in all of these respects and had all served me well, the argument would be stronger.)

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=93#h5p-98

CAUSAL REASONING

Suppose a bunch of people get upset stomachs after a dinner party. Here are the details of what the various people at the dinner party ate:

Foods eaten by people who got sick

- Person A: ham, potato salad, coleslaw
- Person B: ham, rice salad, lettuce salad
- Person C: ham, pasta salad, carrot salad.

Foods eaten by people who didn't get sick

- Person D: chicken, rice salad, coleslaw
- Person E: sausages, pasta salad, lettuce salad
- Person F: bean salad, potato salad, carrot salad

An interactive H5P element has been excluded from this version of the text. You can view it online here:

https://oercollective.caul.edu.au/howtothinkcritically/?p=95#h5p-99

Probably, the ham caused the illness.

Why?

- All those who got sick ate ham
- All those who didn't get sick didn't eat ham
- There's nothing else that was eaten by all and only those who got sick
- Eating bad ham is the kind of thing that we expect to causes upset stomachs – we know roughly how this happens (unlike, for example, something else that the people who got sick might have had in common such as all wearing red shirts).

Note that it could have been true that the ham caused the illness even if it wasn't true that all and only the ham-eaters got sick – eating ham might have raised the probability of someone's getting sick without guaranteeing that they would get sick (some people have cast-iron constitutions).

Causal statements are sentences which say that one thing causes, or doesn't cause, another thing. For example, smoking causes lung cancer, drinking coffee after dinner makes me stay awake, reading logic textbooks after dinner puts me to sleep.

Causal statements are made all the time, both in everyday conversation and in the context of scientific research. We need to know what kinds of effects our actions and other people's actions are likely to have, so that we can decide what we should do in any given situation. Doctors often need to know the causes of diseases in order to know how to treat them. Airlines need to know what caused a particular plane-crash so that they can ensure that the same thing won't cause another one.

Causal *arguments* consist of a causal claim plus the reasons we have for believing that claim. Suppose that American Airlines claims that their plane hit a mountain because the altimeter wasn't working properly and visibility was extremely poor because of low cloud. Their reasons for believing this claim might include records which show that the altimeter was reading fifteen thousand feet just before the plane hit the mountain, when the mountain is nothing like that tall; a tape recording of the pilot's exclamation as he saw the mountain emerge from the fog in front of him, and so on. If you list these reasons as premises and the causal claim as a conclusion, you have a causal argument.

Causal arguments are non-deductive. In the plane case, for example, you can list all the evidence you want and it will still not be 100% certain that you're right about what caused the crash. But it can still be a very good argument.

Consider a more general causal claim: Attending St Peter's Cambridge causes people to get better NCEA results. Suppose we do the stats and it turns out that indeed, the average marks

of students at St Peter's are higher than the average mark for the country at large. Does that provide good reason to believe the causal claim?

No, not on its own. Correlation is not proof of causation. There are other possibilities that you should consider before accepting a causal argument like this one.

1. Coincidence. Might it be pure chance that students at St Peter's did better?
2. Common cause. Perhaps there is some underlying factor which both makes it likely that students will go to St Peter's and makes it likely that they will get good marks: having wealthy parents, perhaps, or having parents who care about their children's education and therefore are likely to both send them to a school with a good reputation and to make sure they do their homework.If you wanted to rule out those alternatives, you should conduct a more complex study. Take a group of students which are the same as your group of St Peter's students in all relevant respects except for which school they go to. Then see if the St Peter's students do better than this control group. If they do, and you've really thought of all other relevant factors, you then have much better reason than before to think your causal claim – going to that school causes you to get better marks – is justified.
3. Opposite direction of causation. I don't think this applies in this case, but sometimes when people infer

causation from correlation, they mistake the cause for the effect. Here is an example. New Hebrides Islanders used to believe that lice caused good health. Why? All the healthy islanders were infested with lice, while sick islanders weren't. In fact, it turned out, the causal connection went the other way. Lice would jump ship when their hosts got a fever, because they don't like high temperatures – so getting sick caused the absence of lice, rather than the absence of lice causing you to get sick. Just from the correlation between health and the presence of lice, you can't infer that the lice cause good health: in this case, the causing went in the other direction.How was the truth about this discovered? I don't know. But it might have been discovered by paying close attention to the order of events. If X doesn't come before Y, then X can't cause Y. (But if X does come before Y, that doesn't necessarily mean X causes Y, of course.)

Having a theory to explain how the causal process in question works is also important. If you discover on independent grounds that lice don't like high temperatures, that gives extra reason to think that illness causes lack of lice rather lack of lice causing illness.

Here is a causal argument:

P. Most people who take mega-doses of Vitamin C when they have a cold recover from their cold within a

week.

_____

C. Mega-doses of Vitamin C cure colds.

We are only justified in believing this conclusion on the basis of this premise if we have considered and ruled out the likely alternatives.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=95#h5p-100

It might be that people naturally recover from colds within a week with or without Vitamin C. This would be easily tested by collected data about the recovery speed of people who don't take megadoses of Vitamin C.

Or perhaps the people who take Vitamin C are people who care about their health and are inclined, when they have a cold, to eat chicken soup and go to bed early, and perhaps those are the factors that cause them to recover quickly, rather than the Vitamin C. The way to test this would be to observe a control group who are exactly like your test group in all relevant respects (diet, sleeping habits, etc) except that they don't take Vitamin C, and compare the two groups.

INFERENCE TO THE BEST EXPLANATION

An inference to the best explanation is when you have some phenomenon or observation that needs explaining, and you conclude that whatever is the best available explanation of that phenomenon is true, just because it is the best explanation of that phenomenon. Sherlock Holmes' argument at the beginning of this chapter is probably meant to be an inference to the best explanation, although he doesn't go through all the steps. He concludes, for example, that Watson has taken up medicine again, from the fact that he has a bulge in the side of his top hat of the sort that would be made by carrying a stethoscope around in your top hat (implicit: only practising doctors carry stethoscopes around with them).

To really spell out the details of the argument, he would need to consider some alternative explanations for the bulge in the top hat, and explain why the "Watson has taken up medicine again" explanation is the best of them. But it seems pretty clear that this is what Holmes has in mind – Watson having taken up medicine again is the best explanation of the bulge in his top hat, so the bulge in his top hat is reason to believe that he has taken up medicine again. Detective story reasoning often takes this form. So does scientific reasoning

– often the reason we have to believe in some type of unobservable entity (such as electrons) is that their existence is part of a good explanation of observable phenomena (e.g. the fact that the lights come on when you flick the switch).

The argument from design (which is an argument for the existence of God) can be construed as an inference to the best explanation. It looks like this.

Observations:

- organisms are complex and intricate
- they are well adapted to their surroundings
- their parts work together to enable the whole organism to function.

Some possible explanations of the observations:

1. God designed organisms to be just the way they are.
2. Organisms evolved by natural selection – no supernatural forces were involved.
3. Organisms evolved by natural selection, but God designed them to do so.
4. God created organisms 6000 years ago in such a way that it would look as though they had been around much longer and had evolved.
5. Organisms came to be the way they are by completely random processes.

5 is not a very good explanation at all. In considering which explanation of some phenomenon is the best explanation, the first thing to take into account is: If the explanation were true, would the observations be surprising, or would they be just what you would expect? Usually when we're looking for an explanation, the thing that we're trying to explain is something surprising – it's *in need* of an explanation. What a good explanation does is render it unsurprising. Explanation 5 doesn't do that – it leaves the complexity and intricacy of organisms surprising.

The other 4 explanations all pass this initial test. Now, to complete the argument, since the conclusion is that God exists, we need reasons for thinking that either 1 or 3 or 4 is a better explanation than 2. There is evidence against 1 (fossils, vestigial organs, etc). However, there is no scientific evidence which would decide between 2, 3 and 4. One reason for preferring 2 might be that it is a *simpler* explanation. One reason for preferring 3 might be that it *explains more*: it can account for the origins of life, whereas 2 cannot. (We are not going to settle this question here – we're just illustrating how inferences to the best explanation work.)

Another example: suppose you observe that, surprisingly, milkmaids don't get smallpox even at a time when smallpox is rife. What's the explanation? Milkmaids get the fairly innocuous illness cowpox and having cowpox gives immunity to smallpox.

Now, the fact that milkmaids don't get smallpox doesn't

conclusively *prove* that cowpox gives immunity to smallpox. There may be other explanations. Perhaps cows have magical powers that protect those who spend much time with them from getting smallpox. Perhaps milkmaids drink more milk than most people, and milk contains some substance which protects people from smallpox.

What makes those less good explanations, or what would show that those were less good explanations?

- Consistency with other accepted theories. What seems like a good explanation to you will partly depend on your background assumptions. I have the background assumption that magic doesn't operate in other areas of the world – I don't need to appeal to magic to explain how I can click on an icon on my computer screen and all your names and student ID numbers come up, for example. You might not share this assumption.
- The results of experimental testing. Do other milk-drinkers have immunity from smallpox? (As it happens, they don't.)

Inferences from evidence to explanations are not deductively valid. It's always possible that the explanation is wrong, in spite of the evidence.

CHAPTER 4 QUESTIONS

Are the following arguments valid, strong, or neither?

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=99#h5p-101

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=99#h5p-102

—

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=99#h5p-103

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=99#h5p-104

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=99#h5p-105

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=99#h5p-106

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=99#h5p-107

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=99#h5p-108

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=99#h5p-109

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=99#h5p-110

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=99#h5p-111

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=99#h5p-112

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=99#h5p-113

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=99#h5p-114

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=99#h5p-115

PART V

CHAPTER 5: SOUNDNESS AND COGENCY

We have so far been talking about the *structure* of arguments: that is, the connection between the premises and the conclusion. However, a good argument needs more than just a strong connection between the premises and the conclusion. In order for it to be an argument we should be persuaded by, it also needs to have premises that are true, or that we have good reason to accept.

A *sound* argument is a valid argument with true premises. A *cogent* argument is a strong argument with true premises. In this chapter we discuss how to assess the truth of premises.

INTRODUCTION: SOUNDNESS AND COGENCY

We've already talked about how to evaluate the *structure* of an argument. This involves thinking about what kind of connection there is between the premises and the conclusion. The terms we have used for talking about the structure of arguments are *valid* and *invalid* (for deductive arguments) and *strong* (for non-deductive arguments).

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=103#h5p-116

There is more to a good argument than validity or strength, however. A valid or strong argument might have premises that you know to be false – in that case you definitely shouldn't believe the conclusion on the basis of the premises, in spite of

the validity or strength of the argument. A good argument, as well as being valid or strong, must also have all true premises. When you are deciding whether you should be persuaded by an argument, you need to think about whether it is *sound* (if it is a deductive argument) or *cogent* (in the case of a non-deductive argument).

Some definitions:

A **sound** argument is an argument that is valid and that has all true premises.

A **cogent** argument is an argument that is strong and that has all true premises.

In this chapter we are going to talk about how to assess whether the premises of an argument are true.

Remember: to be sound, an argument has to be *both* valid *and* have true premises.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=103#h5p-117

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=103#h5p-118

We already know how to assess the validity or strength of an argument. So in order to determine whether an argument is sound or cogent, we need to rule on whether its premises are true.

It is common for people to shy away from making a firm determination of the truth value of a statement. People say things like "it's just an opinion", or "there isn't really one truth". Assessing the truth of statements is a necessary task in argument assessment, and it's important to make an effort to do it.

Is there such a thing as truth?

A statement is true if it describes the world as the world is, and false if it describes things differently from the way the world is. That is a common-sense account of truth, and it is the one we will be using in this course.

Some statements are true in a totally uncontroversial way.

Squares have four sides

is true.

Some statements are equally straightforwardly false.

Squares have three sides.

No matter how forcefully the arguer asserts this, it is simply false.

Some statements are true at some times and false at others. That doesn't stop them from being true or false, in just means that sometimes the world changes, and that will affect the truth of claims about the world. Consider this claim:

Joe Biden is the President of the United States.

At the time of writing, this statement is true. At some time in the future it will be false. This feature of the statement does not stop it being true at the moment.

This way that the truth of claims can change transfers into the soundness or cogency of arguments in which they occur. So,

P1) Joe Biden is the President of the United States.

P2) Joe Biden is a man.

C) The President of the United States is a man.

is a sound argument. At the point at which Joe Biden ceases to be the President of the United States, it will cease to be a sound argument.

Some statements are obviously true. Some, although they are not controversial, you will need to look up, or do some sort of investigation to establish as true or false.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=103#h5p-119

If you needed to look up the answer to the question, that is fine. You're not expected to know everything already.

Sometimes it seems difficult to assess the truth of a statement. Consider

It is raining.

Someone might claim this is true now, but false tomorrow, or true in Hamilton but not in Auckland.

The difficulty here is not created by truth, but by the way the statement is put. Generally, when someone says "it is raining" what they mean is "It is raining here, at the moment". (It would be very strange to say "It is raining", and then say "Oh, I meant it is raining in Thunder Bay, Ontario". That would be to miscommunicate.) We could make such statements more specific (by always including reference to time and place). But it would be tedious to do so, so we don't. It is still either true or false that it is raining here, at the moment.

Sometimes, it is not possible to establish the truth or falsity of a statement. Sometimes, even after research, it remains controversial. When that happens, you may need to suspend

judgement on the truth of the statement, and so also on the soundness or cogency of the argument. It is important, however, to not give up too soon when it comes to assessing truth. You are giving up too soon if you find yourself saying "who really knows?" to just about everything.

There is a philosophical position which involves doubting the truth of nearly everything. It is known as 'scepticism'. Note that even sceptics live their lives *as if* a whole lot of ordinary claims are true. You may adopt the position of scepticism if you choose, but for the purposes of assessing the soundness and cogency of everyday arguments it's important to take a more everyday approach to truth.

Some propositions are more difficult to assess than others. It is important to not simply give up and throw your hands in the air in horror. Pause, and think carefully.

Consider the statement

> Most of the people reading this textbook are enrolled in a course.

Is this true?

When considering its truth value you might start by noting that most people in the world are not enrolled in any sort of course. So if it's true, it needs to be something special about people who read this textbook.

We can also note that it is not the case that all people who read the textbook are enrolled in a course. I have read it, for instance, but I'm not enrolled in any course.

If you're reading this textbook because it's assigned for the

course, then you know that it is true that *at least some* people reading it are enrolled in a course. But how do you work out if it is true that *most* people who read it are enrolled in a course?

The answer is that you should think carefully about whether it is likely to be true. It is *possible* that this textbook has secretly become some sort of bestseller, and that lots people are reading it for fun. But it's not very likely. Generally speaking, people read textbooks to help them with some particular course. The number of people who read textbooks for recreational purposes is fairly small, and this textbook is less accessible than many, because you can't walk into an ordinary bookshop and purchase it. So, having weighed all that up, it seems to me that the statement is likely to be true.

Determining truth is not always easy, but it is possible, and it is important.

Cogent arguments and false conclusions

With a deductively valid argument, the truth of the premises guarantees the truth of the conclusion. That means that a valid argument with true premises must have a true conclusion. This makes sound arguments particularly useful.

With non-deductive arguments, the relationship between the premises and the conclusion is different. In a strong argument, the truth of the premises makes the conclusion likely, but stops short of guaranteeing it. One consequence of

this is that it is possible to have a cogent argument with a false conclusion.

Consider the following argument:

> P1. Nearly all of the presidents of the United States have been white men.
>
> ---
>
> [Probably] C. The 44th president of the United States was a white man.

This is a strong argument: the truth of the premise does not guarantee the truth of the conclusion, but it does make it very likely. There is no problem with the form of the argument.

The premise is true. (There have been 46 US presidents, and 45 of them have been white men.)

Since we have a strong argument with true premises, this argument is cogent.

Note, however, that the 44th president of the United States was Barack Obama. He is not a white man. So this is a cogent argument with a false conclusion.

We can make this even more explicit:

> P1. 45 of the 46 presidents of the United States have been white men.
>
> P2. Barack Obama was the 44th president of the United States.
>
> ---
>
> [Probably] C. Barack Obama was a white man.

This argument is still a cogent argument.

This is merely a consequence of the way cogency has been

defined. Any strong argument with true premises will be a cogent argument *even if the conclusion is false*.

This means that there is an extra step when assessing cogent arguments than there is with sound arguments. With a sound argument, the truth of the conclusion is guaranteed, and does not need to be independently assessed. With a cogent argument, once the argument has been established as cogent it's a good idea to consider whether you have any additional reason for thinking that the conclusion is false. It isn't common, but sometimes there will be some additional information you possess which shows you that the conclusion is false. This won't stop the argument from being cogent (because that is determined by the definition of cogency), but it does deserve to be commented on. It is a relevant piece of assessment for cogent arguments.

ASSESSING NORMATIVE CLAIMS

Arguments often have normative conclusions. That is, they make claims about what we should or shouldn't do, or about what ought to be the case. To be valid or forceful, an argument with a normative conclusion must have a normative premise (a premise with a "should" or an "ought" in it).

To assess the soundness of such arguments we need to assess the truth of normative claims.

Moral claims are one common type of normative claim. Some people think that it is not possible to assign truth values to moral claims. There are severe disadvantages to such a view. If moral claims cannot be true or false, then there are no actions which are not permissible. Most people don't want to accept that. So, if there are no moral truths, then you cannot say that

> Slavery is wrong

is true. Nor could you say that

> It is okay to torture babies for fun

is false.

In fact, if moral claims cannot be true or false, we would never be able to assess an argument which included a moral claim as a good argument. This is unsatisfactory. We need to

be able to assess moral arguments, and to establish that some actions are impermissible, and some are good, and so forth.

In this course, we accept that moral claims, and other normative claims, can be true or false.

That does not mean that it is always easy to know the truth value of a normative claim. Like other propositions some are easier to evaluate than others.

Suppose someone gives the following argument:

> Smoking can cause cancer, and it's bad for your lungs. This means smoking is bad for your health, and therefore you should stop smoking.

When we reconstruct this argument, we'll end up with a connecting premise that says something like

> If smoking is bad for your health, then you should stop smoking.

Is this true?

Your first reaction might be that it's true. But we should think carefully. Note that the more general claim "If something is bad for your health then you should not do it" is *not* true. There are exceptions. For example, fire-fighting is bad for your health. Fire-fighters regularly suffer burns, and smoke inhalation, and it's a stressful job. All these things are bad for the health of fire-fighters. But that doesn't mean we shouldn't have fire-fighters. Or, being a defensive tackle on an American Football team is bad for one's health. They have a high risk of serious concussion and other injuries. But it's not clear that it

follows that all defensive tackles should immediately quit their jobs.

So the statement "if something is bad for your health then you should not do it" is not true. Is the smoking-specific version of it true?

When you ask someone why they smoke, a variety of reasons are given. It brings pleasure, it's a social activity (people who smoke often do so with other people), it's a useful way to unwind, and so forth. There's not a greater benefit to society the way there is with fire-fighting. Nor are there the financial rewards and the personal satisfaction of being a good defensive end. If the people who currently smoke ceased smoking, there wouldn't be the drawback for society that there would be if fire-fighters all quit their jobs. Nor would any smokers lose their livelihood as a result.

We can also note that there are other ways to achieve the benefits of smoking: there are other ways to relax, and other ways to be social. Could you vape with your friends instead? (Whether you can may depend a bit on your friends.) Vaping is healthier than smoking, but can be done in many of the same circumstances, and delivers nicotine, so it can accomplish many (though presumably not all) of the same things as smoking.

We can also note that the claim made in "If smoking is bad for your health, then you should stop smoking" is not a claim about what you *will* do, but about what you *should* do. Given that, and that vaping, which is healthier, can achieve many of

the same things, I'm willing to accept that "If smoking is bad for your health, then you should stop smoking" is true. I only accept it, however, in the smoking case. I still hold that "If something is unhealthy then you shouldn't do it" is false.

This gives you some idea of what should be considered when assessing the truth of a statement.

We can now insert the statement back into the argument.

P1) Smoking significantly increases your risk of getting cancer.

P2) Smoking is harmful to your lungs.

P3) Anything that significantly increases your risk of getting cancer and is harmful to your lungs is bad for your health.

C1) Smoking is bad for your health.

P4) If smoking is bad for your health, then you should stop smoking.

C) You should stop smoking.

P4 was the most controversial premise here. I haven't specifically evaluated the others, but they are true. Given that P4 is also true, this argument is sound.

ASSESSING AUTHORITIES

Sometimes we decide that a premise is true not by thinking it through (as we did above with the smoking example) but by accepting it because someone else says it's true. In some cases that's a perfectly reasonable thing to do: in some cases it isn't.

Joe at the bus-stop says that interest rates are about to go up. Should I believe him? Not straight away. There are two lines of questioning I could engage in. I could ask for his credentials, and if he's a banker I might think that was reason enough to accept what he says. Or, I could ask for his evidence – what reasons does he have for thinking that interest rates are about to go up? – and then evaluate his reasons for myself.

If I am competent to evaluate his reasons, then I should probably do the second of those things (though of course sometimes there just isn't time to do this properly). But if you're doing the first thing – accepting what someone says because they're an expert, without asking what their reasons are for believing what they're saying – then you should check whether or not the following conditions apply:

1. The person must be a genuine authority in the area in question.

2. There must be substantial agreement between authorities.
3. The authority must be testifying honestly.

If any one of these conditions is not met, the appeal to authority is fallacious.

1) The person must be an authority in the relevant area. Sportsmen endorsing sports shoes is one thing: sportsmen endorsing insurance companies is another. If an All Black tells me Nike are the best, I might reasonably (setting aside for the moment the question of how much he's been paid to say so) think that he knows what he's talking about: presumably All Blacks do have considerable experience with sports shoes. On the other hand, if the All Black tells me that Tower is the best insurance company, or touts the nutritional benefits of Weetbix, I should think twice. Is there any reason to think that what an All Black says about insurance is more likely to be true than what Joe at the bus-stop says about insurance? Is there any reason to think an All Black is an expert on nutrition? Probably not.

2) There must be agreement amongst experts in the area in question. If the issue is one that authorities almost all agree on, like which planet in our solar system has the greatest mass, then the appeal to authority may be legitimate. If there is no agreement even amongst the experts, then the appeal to authority is not legitimate.For example, there is no agreement amongst experts on horse-racing about who is the greatest

racehorse of all time. Suppose that the ten top horse-racing experts back ten different horses. Under those conditions, if I conclude because expert Z says so that Phar Lap is the greatest racehorse of all time, I don't really have very good reason at all to believe that conclusion.

3) If there is good reason to doubt the honesty of an expert, then appealing to what that expert says is not a legitimate appeal to authority. For example, in a trial, if you discover that the expert called by the prosecution is being paid to testify in a certain way, that discredits their testimony.

It is particularly difficult and time-consuming finding out whether or not all or most experts are in agreement. In the situation in which you're on a jury and you're trying to evaluate the testimony of expert witnesses, you can probably assume that if experts disagree the attorneys will tell you this. If the defence's expert witness is the only one who believes the claim that they're making about DNA matching, you can rely on the prosecution, if they're halfway competent, to point this out.

In everyday life, there isn't time to check everything. We're bombarded with information from all sides, and while we shouldn't believe every word of it, we can't check the credentials of every person (or every magazine or newspaper) either. If you have a normally reliable source, you will probably be inclined to believe what they say unless it sounds, based on your own experience, highly implausible. That's okay. Most of

it doesn't matter very much. If it does matter, you are, and should be, more careful in your checks.

SOME RELEVANT FALLACIES

In this section we consider the following fallacies:

- Fallacious appeal to authority
- Fallacy of majority belief
- Appeal to tradition

Fallacious appeal to authority

If someone says you should believe something because an expert says it is true, and any of the three conditions mentioned above for a legitimate appeal to authority are not met, then you have a *fallacious appeal to authority*.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=109#h5p-120

The fallacy of majority belief

A fallacy of majority belief occurs when someone uses the fact that most people believe something as justification for why that thing is true.

The reason this is fallacious is that it is perfectly possible for most people to believe something that is false. So, when most people believed that the world was flat, that didn't make it flat. It is just that most people believed something which was false. And when most people believed that slavery was justifiable, that didn't show that slavery was right. It just means that most people were wrong.

It is much harder, of course, to see that most people are wrong when sitting within that consensus. What we can do is acknowledge that this general pattern of argument fails. Any argument that appeals to widespread belief as its justification for a claim commits a fallacy of majority belief.

Consider this argument:

> P1) Most people in my hometown agree that the earth is less than 10,000 years old.
>
> ---
>
> C) The earth is less than 10,000 years old.

In order to make such an argument valid, the connecting premise would need to say something like

> P2) If most people is my hometown agree on a thing, then that thing is true.

It's that connecting premise, often omitted from the argument, which is false.

That example is an easy one to spot. Many examples of the majority belief fallacy are more subtle. Consider this argument.

> A recent poll showed that 77% of New Zealanders support the legalisation of the medical use of cannabis. Medical cannabis should be legalised.

Reconstructed, the argument looks like this:

> P1) 77% of New Zealanders support the legalisation of medical cannabis.
>
> P2) If 77% of New Zealanders support something, then New Zealand should do that thing.
>
> ---
>
> C) New Zealand should legalise medical use of cannabis.

We have a tendency to support P2: we might feel, for instance, that democracy requires that we do what most people want. But P2 is false. There are all sorts of thing that would be a bad idea to do even if 77% of New Zealanders thought it was a good idea. You might be able to get 77% of New Zealanders to agree that they should all be given a million dollars. But doing so would generate real problems.

Often, of course, the majority of people will believe something that is true. What the fallacy of majority belief brings out is that the fact that most people believe something

is not, by itself, a reason to accept that something is true. You should look for other reasons.

Appeal to tradition

Appeals to tradition are like appeals to majority belief except that rather than appealing to the *number* of people who believe something as a reason to believe it, they appeal to *how long* people have believed it.

Here's an example:

> *People have believed that there is a God or Gods for thousands of years – so it must be true that there is a God or Gods.*

You can probably immediately see that the objections to arguments from majority belief apply here, too. The mere fact that some belief has been around for a long time is not enough reason for *you* to adopt it. Instead, consider what reasons those thousands of years of humans had for believing in God, and see whether or not they were good reasons. If so, then it's those substantial *reasons* that are doing the work of supporting your belief in God, not the fact that people have been believing in God for thousands of years.

An interactive H5P element has been excluded from this version of the text. You

can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=109#h5p-121

CHAPTER 5 QUESTIONS

Here are some questions to practice on.

For each of these statements, indicate whether the statement is true or false.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=111#h5p-122

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=111#h5p-123

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=111#h5p-124

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=111#h5p-125

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=111#h5p-126

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=111#h5p-127

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=111#h5p-128

For the following questions you need to determine whether they are sound or unsound.

Remember, to be sound an argument has to be valid, *and* have true premises.

An interactive H5P element has been excluded from this version of the text. You can view it online here:

https://oercollective.caul.edu.au/howtothinkcritically/?p=111#h5p-129

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=111#h5p-130

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=111#h5p-131

An interactive H5P element has been excluded from this version of the text. You

can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=111#h5p-132

For the following questions you need to determine whether the arguments are cogent or not cogent.

Remember, to be cogent an argument has to be strong, *and* have true premises.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=111#h5p-133

An interactive H5P element has been excluded from this version of the text. You can view it online here:

https://oercollective.caul.edu.au/howtothinkcritically/?p=111#h5p-134

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=111#h5p-135

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=111#h5p-136

PART VI
CHAPTER 6: FAIR DEALING

This chapter is about various ways in which people can fail to deal fairly with other people's arguments and ideas. For example, if you write a comment on a blog post and others completely misrepresent what you've said and slag off at you for things that you didn't in fact say, they are not dealing with you fairly.

It's annoying to be on the receiving end of unfair dealing. We will give you some tools for recognising and responding to it.

We also encourage you to deal fairly with other people's arguments. The various kinds of unfair moves we will be discussing might sometimes "win" an argumentative discussion, if the person you are arguing with is not alert to them. But you should not be solely interested in winning – where "winning" means having the other person give up and either agree that you are right or just not bother to engage with you anymore. It's much better to persuade people by clearly and straightforwardly presenting genuinely good reasons in support of your view than by brow-beating them or confusing them.

Unfair dealing attempts to shut down other people's attempts to say what they think and have their arguments heard. There are two different kinds of reasons why this is bad. The most obvious one is that it's bad for the people whose arguments you are suppressing or dismissing – put yourself in their shoes, and you will see how frustrating that would be, and how much it might make you feel as though you are not being treated as someone worth listening to. This is an ethical reason to be a fair dealer.

The other kind of reason has to do with what kind of approach to arguing is most likely to get you, and the people you are talking to, closer to the truth. If you are genuinely interested in finding out what there is good reason to believe, or deciding what really is the best thing to do in some situation, then you should encourage others to say what they think and why they think it, open-mindedly consider what they say, modify your own view if they have presented good reasons why you should, and clearly and calmly put forward your own arguments so that they can point out any problems with them. This is fair dealing.

Sections:

- The Burden of Proof and Arguments from Ignorance
- Chapter 6 questions

THE PRINCIPLE OF CHARITY

Sometimes when you are faced with an argument, you can see immediately what it means and what the arguer's intentions are, but sometimes not. Often you have to do a bit of interpretation.

For example, sometimes one or more premises in an argument are left unstated. This doesn't necessarily mean the argument is a bad one – it may be that when you spell out the implicit premises, they are so obviously true that there is no problem. That is in fact one reason why you might leave premises unstated – they're so obvious that they don't seem worth mentioning.

The problematic case is when the unstated premises are implausible. Sometimes sneaky arguers deliberately leave crucial premises unstated because those premises are completely implausible, and hope that the people they are trying to persuade don't notice that the argument relies on an implausible unstated premise.

Sometimes even the conclusion of an argument is not explicitly stated. In advertising, for example, the conclusion tends to be "You should buy this product," but often that's not explicitly said – you're told the reasons why you should

buy it (the premises) but left to figure out the conclusion from the context. Likewise in politics. In an election year you hear many arguments whose conclusion, not always explicitly stated, is "You should vote for me!"

Issues of fair dealing arise both when you're trying to persuade someone by giving an argument (you should avoid sneakily leaving controversial premises unstated) and when you're evaluating someone else's argument (if the argument is not completely stated or not completely clear, you should be charitable in your reconstruction of it).

Here's an example of an argument which relies on an unstated premise or assumption. See if you can figure out what premise needs to be added to get from the premises of this argument to the conclusion. Hint: it's a premise that's so obviously true that you might have trouble even noticing its absence.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=115#h5p-137

In the "microwave dinners" example, the unstated premise is one that we would probably all accept.

But consider this argument:

> *Murder is the killing of an innocent human being. So abortion is murder, because abortion is the killing of a foetus, and you can't get much more innocent than a foetus.*

This also has an unstated premise: it relies on, but does not explicitly state, the premise: *Foetuses are human beings*.

In this case the omission may be less benign. The unstated premise is controversial, and the arguer may be hoping that if they don't state it, their opponent won't notice that the argument relies on it.

When you are reconstructing an incomplete argument, you should apply the **principle of charity:** where you have a choice about how to reconstruct the argument, you should reconstruct it in such a way as to make it the best argument it can be. That is, if you can, you should reconstruct it as an argument with plausible premises and a strong connection between premises and conclusion.

Here's an example. Consider this incomplete argument:

> *Abortion is deliberately killing an innocent human being, so obviously it's morally wrong.*

There are a number of different premises you could add to get from the premise "Abortion is deliberately killing an innocent human being" to the conclusion "Abortion is morally wrong." In cases like this in which you have a choice, how do you decide what premise to add?

You should:

a) Use whatever evidence you can get about the arguer's intentions from the stated premises, conclusion and context.

b) Apply the Principle of Charity. This principle says that when faced with an argument which needs reconstructing or augmenting, you should reconstruct it in as as charitable a way as possible. If you can avoid it, you shouldn't add premises that are obviously false – you should add the most plausible premise that will do the job. And you should add premises which help to link the stated premises to the conclusion.

Why be so nice?

Well, if you're in a debate with someone and you attack a version of their argument which isn't as strong as it could be, be, the person will just say: "That wasn't what I meant. You're not attacking my actual argument, you're caricaturing my argument." So you won't have got anywhere.

Consider the example above. Someone says to you: "Abortion is deliberately killing an innocent human being, so obviously it's morally wrong."

There's a premise missing – something that connects the killing with the wrongness – and there are a number of different premises that would do the job.

1. Killing is morally wrong. (No – many people think it's okay to kill animals for food, and almost everyone thinks it's okay to kill plants for food.) So let's try:
2. Killing humans is morally wrong. (No – we think it's okay to kill in self-defense.) So let's try:

3. Killing innocent humans is morally wrong. (No- it's not morally wrong to kill someone by unavoidable accident. If a kid runs out in front of my car and I hit them because I couldn't stop in time despite my best efforts, that's extremely regrettable, but it doesn't seem as though I did something morally wrong.) So let's try:
4. Deliberately killing innocent humans is morally wrong.

If you reconstruct your opponent's argument by adding 1 above, and then attack 1, the person will brush you off by denying that 1 was what they intended to say. It's tempting to do this, because it's much easier to find counterexamples to 1 or 2 than to 4. But you won't win the argument that way, unless your opponent is half asleep. If you attribute to them the more plausible unstated premise 4, which is probably what they intended, then you can have a productive debate about the issue and your opponent won't be able to side-step your objections in the same way. Is 4 true? Is it true that abortion is the deliberate killing of an innocent human being? Having reconstructed the argument in the most charitable way you can, it's now possible to evaluate it – to consider whether the premises are true, and whether, if true, they provide good reasons to believe the conclusion.

FAILING TO ENGAGE: AD HOMINEM, STRAWMANNING, AND WHATABOUTERY

When someone gives an argument, the proper way to respond is to assess their argument. That is, you should listen to what they say, and consider whether their argument is sound or cogent. There are a number of common ways in which people can fail to engage in with an argument that someone else has given.

The section discusses the following:

- Ad Hominem Arguments
- Strawman Arguments
- Whataboutery

Ad hominem arguments

Here is a (slightly adapted) argument which appeared in the NZ Herald.

Don't terrorise legal drinkers

> **by Mike Douglas** [Mike Douglas is the CEO of a major brewing company.]
>
> *The current campaign against combining drinking with driving is terrorising law-abiding people. Many law-abiding people are cutting their alcohol consumption because they are afraid of being caught by random breath testing. A 2014 study of more than 4000 people showed that more than a third were drinking less or had stopped altogether compared to the previous year. The prime reason they gave was the risk of being caught by random breath testing.*
>
> *But research shows that the average drink-driver in a fatal accident has an average blood alcohol level of more than twice the legal limit. The current campaign against drinking and driving is failing to achieve what should be our top priority; getting the heavy and hardened drinkers off the road.*

Now suppose I say, in response to this:

> *Douglas has a vested interest in getting people to believe this conclusion, since he is the CEO of a major brewing company. He is biased, so I reject his argument.*

I am committing the ad hominem fallacy: trying to discredit an argument by attacking the source of the argument. Why is this a fallacy? Well, if an argument is a good argument, it doesn't become less good because of any facts about the arguer. If we're given a valid argument with plausible premises which causes us to believe its conclusion, and then we discover that the arguer is evil, we shouldn't change our opinion of the argument. Likewise if we discover that the arguer is incredibly

stupid. An argument is an argument, and it should be assessed on its merits. Douglas has provided reasons in support of his conclusion, and I should evaluate those reasons. Are his premises true? If true, do they support his conclusion? I can't (or at least, I shouldn't) get out of considering these questions by writing off the argument because of facts about Douglas.

> **Ad hominem fallacy:** when someone tries to discredit *an argument* by attacking the source of the argument.

However, attempting to discredit someone's *testimony* by attacking the person testifying can be legitimate. If I testify that I saw one of my colleagues stealing stationery from the philosophy secretary's office, then facts about my reliability as a witness may be highly relevant. If I'm a compulsive liar, that gives you reason to doubt my testimony. If I'm known to have stolen stationery myself in the past and tried to pin the blame on others, that's relevant too.

But now suppose I give an argument:

> *The stationery was stolen between the hours of 5pm on Monday and 8am on Tuesday. The security system records that the only person to enter the building during those hours used Tracy's pin number when he or she turned off the alarm system. No-one else knew Tracy's pin number. So it is extremely likely that Tracy was the thief.*

Attacking my character was legitimate as a way of undermining my testimony, but it doesn't undermine my argument. Now that I've given an argument, the argument

should be evaluated on its merits. We know how to evaluate an argument on its merits from earlier chapters: you have to think about whether the premises are true, and you have to think about how well they are connected to the conclusion. Writing my argument off because I'm known to be a dishonest person fails to do either of those things.

Of course, some ad hominem attacks are bad arguments even if they're not ad hominem fallacies as such. Claims made about a person may be irrelevant to whatever is the point at issue, or they may be false. For example, if a jury member argues that the defendant must be guilty because they have beady eyes and a weak chin, this is not a good argument, because these sort of facts about a person's appearance are irrelevant to the question of whether or not the person is guilty of some crime.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=117#h5p-138

Strawman arguments

A strawman argument is when someone distorts or caricatures

an opponent's arguments or views, and then attacks the weakened version rather than the real argument.

Here's an example. In the previous section we saw an argument by Mark Douglas about a campaign to reduce drinking and driving. Now suppose someone responds to that argument as follows:

> *Mark Douglas argues that we shouldn't have laws against driving while intoxicated. But that's absurd. It's well-known that alcohol impairs judgement and affects coordination. We need to keep drunk people off the roads.*

The arguer is "strawmanning" Mark Douglas. If you look back at Douglas's argument, he was not arguing that we shouldn't have laws against driving while intoxicated – merely that the current campaign isn't a good way of responding to the problem of people driving while intoxicated. So the arguer here is caricaturing Douglas's argument and then attacking the caricatured version, rather than addressing Douglas's actual argument on its merits.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=117#h5p-139

Whataboutery

Also known as "whataboutism", whataboutery is a common technique in internet discussions and social media.

The basic technique of whataboutery is to respond to someone's point or argument about X by saying "What about Y?", where Y is not currently being directly discussed. The basic function of whataboutery is to derail the current discussion. The person who gave the original argument finds themselves defending their choice of topic (rather than the argument they gave within the topic), or dealing with a matter quite different from the one they wanted to discuss. The person who has performed the whataboutery move has deflected or misdirected the discussion. Whataboutery can thus take the form of the more specific techniques above: sometimes it turns into an attack on the original arguer, by accusing them of not caring about some other (important) thing. Sometimes it accuses the original arguer of hypocrisy: if they had a consistent position, they would care about this other thing. Sometimes it misrepresents the original arguer's position, and so they find themselves scrambling to explain their true position instead of their original argument being addressed. Whataboutery is always a poor argument move, because it always fails to respond to the the argument which is in fact given.

Consider the way that the "Black lives matter" campaign was responded to by people saying "All lives matter". This is

a whataboutery response: it's the equivalent of saying "What about non-black lives?".

Such a response is a good example of two types of moves at once. First, this is strawmanning the original position. Someone who says "Black lives matter" is not claiming that "Only black lives matter". Someone who responds "All lives matter" is understanding "Black lives matter" as "If a life matters, then it's a black life" when the correct position is "If a life is black, then that life matters". Misrepresentation of a position is attacking a straw man. Second, there's an ad hominem move involved. The whataboutery move here was often accompanied by saying that the "Black lives matters" campaigners were themselves racist, and so inconsistent, and therefore shouldn't be listened to.

Whataboutery is thus not only a failure to engage in the argument, it is also often an aggressive or abusive move. It is a practice which should be avoided.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=117#h5p-140

QUESTION-BEGGING ARGUMENTS

When someone uses their conclusion as one of their premises, they have a valid argument. If all the premises are true, then of course the conclusion must be true too, if the conclusion is just a restatement of one of the premises.

The simplest sort of question-begging argument is an argument where the conclusion simply repeats the premise. So,

P1) Joe Biden is the president of the US.

C) Joe Biden is the president of the US.

Such an argument is valid, and the premise is true (at the time of writing), so it is sound argument. It is, however, totally useless. The general purpose of arguments is to persuade people of things. This argument cannot be used to persuade anyone of anything, because no one would accept the premise unless they already accepted the conclusion. The failure of question-begging arguments lies not in some technical failure in structure, but in the way they fail to serve the purpose of arguments.

Here's an only slightly more complex example:

The sky is blue, and the grass is green, therefore the sky is blue.

This is valid. And its premises are true. However, it is not a good argument. You're not going to convince anyone who is in doubt about the blueness of the sky by presenting this argument, because for someone to be convinced by an argument they have to accept the premises and then see that the premises give reason to believe the conclusion. But no-one who doesn't believe the conclusion already will accept the premises of this argument, since the conclusion *is* one of the premises. This is also a question-begging argument.

No-one would mistake the argument about the blueness of the sky for a good argument. Real-life question-begging tends to be a bit more subtle, and can be a way of trying to disguise the fact that you don't have a good argument for the thing that you are trying to get someone to believe.

Here is an example:

> *Abortion is wrong, for it is always wrong to voluntarily and purposefully destroy a living and growing human foetus.*

In this argument, the conclusion ("Abortion is wrong") is just a restatement of the premise in different words. "The voluntary and purposeful destruction of a living and growing human foetus" is just a more emotive way of saying "abortion." So unless you already believed the conclusion, you would never accept the premise.

If you think about what the point of giving an argument is, you will see why this is problematic. Generally, when we give an argument, we are trying to persuade someone of a

conclusion that they don't already believe. We do it by providing premises that they *are* inclined to believe, and then showing them that if they accept those premises, then really they ought to accept the conclusion as well. A question-begging argument can't be used in this way, because there is no one in the target position – no one who believes the premises but doesn't believe the conclusion – since the conclusion *is* one of the premises.

Another, slightly more subtle, example:

> *God exists. We know this to be true, since the Bible tells us that God exists. And we know that what the Bible tells us is true, since the Bible is the word of God.*

Putting this in standard form:

> P1. The Bible is the word of God
>
> P2. If the Bible is the word of God, then what the Bible tells us is true.
>
> ---
>
> C1. What the Bible tells us is true.
>
> P3. The Bible tells us that God exists.
>
> ---
>
> C. God exists.

This one is more complicated. The premises don't explicitly include the claim that God exists. But they do presuppose it. If you didn't already think there was a God (which is to say, if you didn't already believe the conclusion), you wouldn't accept the premise "The Bible is the word of God." The argument is question-begging – it assumes the very thing that it is trying

to prove. Question-begging arguments are sometimes called *circular* arguments, and you can see why in this example – it argues in a circle, starting from the existence of God and ending up back where it started, at the existence of God.

Question-begging arguments are a reminder that soundness and cogency, while useful, are not the only things that we should consider when assessing arguments.

THE BURDEN OF PROOF AND ARGUMENTS FROM IGNORANCE

Sometimes people engaged in a disagreement try to get out of providing reasons in support of their view by claiming that their opponent has the burden of proof. That is, they claim that their own position is the default position and doesn't need to be supported by arguments, and that it is up to their opponent to prove them wrong.

Here's an example. Suppose I'm having a dispute with a pesticide manufacturer about whether using *NoMoreBugs* on your tomato plants is bad for your health. And suppose the manufacturer says "*NoMoreBugs* is safe. You have no evidence that it's not! We're going to keep selling *NoMoreBugs*."

Hang on a minute. Why is it up to me to show that *NoMoreBugs* is unsafe, rather than up to the manufacturer to provide evidence that it *is* safe? The manufacturer's argument is pretty much:

P1. You haven't proved that *NoMoreBugs* is unsafe.

C. *NoMoreBugs* is safe.

When someone argues that because there's no proof that

their claim is false, therefore it's true, they are committing a fallacy called an **Argument from Ignorance** – arguing that since we don't know that the claim is false, it must therefore be true.

There are all kinds of claims that are extremely likely to be false even though we can't *prove* them false.The fact that we can't prove something false is not enough reason for us to believe it. We can't prove that it's false that 50% of the people in this class are extraterrestrials pretending to be humans. That doesn't mean that we have reason to believe that 50% of us *are* extra-terrestrials pretending to be humans.

Sometimes arguments from ignorance are a bit more subtle. Suppose a tabloid newspaper publishes some scandalous allegations – "Stop Press! Winston Involved In Threesome With Jacinda and Clarke!" And then suppose the next issue contains a piece about how Winston and Jacinda and Clark have failed to deny these allegations, which concludes that the allegations must therefore be true. This is an appeal to ignorance. The fact that no one has produced evidence that the claim is false is not enough reason to believe it's true. There are lots of reasons why Winston and Jacinda and Clarke might not deny the claim even if it's false. For example, "NZ PM and Deputy PM deny allegations..." might make a headline in a reputable mainstream newspaper, thus distributing the unsubstantiated allegations to a wider audience, and that might be something the people involved would rather avoid.

Arguments from ignorance involve illegitimate attempts to *shift the burden of proof.*

In a courtroom, the burden of proof rests with the prosecution. The defendant is considered innocent until he or she is proven guilty. It's up to the prosecution to prove the defendant guilty. The defense does not have to prove the defendant innocent – they only need to show that there are holes in the prosecution's arguments; that there is a reasonable doubt about whether the defendant is guilty.

In the courtroom situation, we know who has the burden of proof – it's built into our justice system. But in other cases, it's not so clear, and there's no one story which applies across the board about what makes it so that the burden of proof is on one person rather than another.

If I'm claiming something which conflicts with what is generally accepted, then the burden of proof is on me – I shouldn't expect my unorthodox view to be accepted by others unless I provide evidence or arguments for it. When Copernicus claimed that the earth goes round the sun rather than vice versa, he needed to supply evidence – the burden of proof was on him – because the traditional view was that the sun goes round the earth. But if I now get into a dispute with someone about this question – I claim the earth goes round the sun, they claim the sun goes round the earth – then the burden of proof is on them – nowadays, the general view is that the earth goes round the sun.

There are other considerations which are relevant to who

has the burden of proof. One of them is safety. On the whole, if people's safety is in question, the burden of proof is on the people arguing that the thing is safe to prove that it is. Why? The consequences of being too careful are much less bad than the consequences of not being careful enough.

But note that this can open the way to malicious, unsubstantiated objections. If a particular medication has only been in use for 10 years, there is no way to conclusively prove that it has no long-term side-effects, and it would be unreasonable to require a drug company to prove that it has none if there is no particular reason to suppose that it has.

Another consideration is whether your view conflicts with a consensus amongst authorities. Suppose that all the most-respected theorists in the field of education agree that exams are not a good way of testing students' knowledge and abilities – supposed the received view in education is that exam results mostly reflect how good the student's exam-sitting technique is, and that internal assessment is much more useful. Then, if I want my course to be assessed entirely by exam because I think that actually making people write down a whole lot of information under time pressure actually is the best way to tell how much they've learned, it's reasonable that the burden of proof should be on me – if I'm flying in the face of what all the educational authorities agree on, I should make a case for my view – I should present arguments.

So there are these are three kinds of cases when the burden of proof is on you:

1. When you are making a claim that conflicts with common wisdom,
2. when you are claiming that some product or practice is safe, and
3. when you are making a claim that conflicts with expert consensus.

Almost always, in fact, you should be willing to provide evidence in support of your claims – why should anyone believe you, about any controversial matter, if you don't provide reasons in support of your view? But these are three kinds of case in which it's *particularly* important to be able to provide reasons in support of your view.

Note that in saying that the burden of proof is on you if what you are claiming conflicts with what is commonsensically believed, I am *not* saying that the fact that most people believe something is a good enough reason for you to believe it too. It never is. For example, it's traditional to think that we don't need to be nearly as considerate of the needs of animals as we should be of the needs of humans, and most people still believe that's true. But I shouldn't believe that, and you shouldn't believe that, just because most people do or because it's a belief with a long and illustrious pedigree. That most people believe it doesn't show that it's true. It might, however, put the burden of proof on the person who disagrees with it – if they want to change people's minds, they will need to provide a compelling argument.

CHAPTER 6 QUESTIONS

For the first set of questions it will help to review other fallacies taught earlier in this book. As well as those we have introduced in Chapter 6, there were structural fallacies in Chapter 2, and fallacies at the end of Chapter 5.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=123#h5p-141

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=123#h5p-142

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=123#h5p-143

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=123#h5p-144

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=123#h5p-145

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=123#h5p-146

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=123#h5p-147

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=123#h5p-148

The following questions review some of the basic terminology of the book.

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=123#h5p-149

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=123#h5p-150

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=123#h5p-151

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=123#h5p-152

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=123#h5p-153

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=123#h5p-154

An interactive H5P element has been excluded from this version of the text. You can view it online here:
https://oercollective.caul.edu.au/howtothinkcritically/?p=123#h5p-155

APPENDIX: YOUTUBE TRANSCRIPTS

Chapter 1

Standard form

In this short video we're going to take you through the basics of putting arguments into standard form. So here's our argument: "It's not the case that everything in the universe is physical unless consciousness is physical. Consciousness is not physical. Therefore not everything in the universe is physical." Our first task is to identify the conclusion. In this argument this is very easy, because we have this indicator word "therefore" which shows us where the conclusion is. The conclusion is "not everything in the universe is physical". So we can label that as our conclusion. Then we look for the premises. Again in this simple argument they are very easy to find. The first premise is "it is not the case that everything in the universe is physical unless consciousness is physical". We can call that premise 1. And the second premise is "consciousness is not physical". We'll call that premise 2. Now when we put the argument into standard form it's often easiest to write the conclusion down first: "Not everything in the universe is physical". In standard form the conclusion is always

separated from the premises which support it using an inference bar. You can just draw a line using a draw function. Or if you're writing your answer in a Word document you can underline the final premise. Then we write in the premises: "It's not the case that everything in the universe is physical unless consciousness is physical" and "consciousness is not physical". We also, when we put arguments the standard form, always label the premises and the conclusion. This makes it easy to refer to them later. So we call the first premise P1, and the second premise P2, and the conclusion C. The argument is now in standard form. What it's important to note about what we've done, is that these features – P1, P2, the inference bar, and the conclusion – are what we call the 'frame' of the standard form rendition. Our standard form is now complete. In later videos we'll look at some more complex arguments and put them into standard form also.

Chapter 2

Modus ponens

It's often easier to identify an argument form if you take the argument and replace the particular statements that occur in the argument with letters. So this video is going to take you through an easy example of how to do this. It's going to use the example of a Modus ponens argument. So here's our basic argument: "If cucumbers are fruit then so are tomatoes. Cucumbers are fruit. Therefore tomatoes are fruit." Now the

first thing you might notice about this argument is that this phrase "so are tomatoes" in P1 actually means the same thing as "tomatoes are fruit". So let's replace that with "tomatoes are fruit". This will make it easier to see how our form works in this argument. P1 is made up of this "if" claim, and then part of its statement is "cucumbers are fruit". We then have a "then" phrase, and the consequent there is "tomatoes are fruit". Now I've used green for "cucumbers are fruit" and red for "tomatoes are fruit" to help us see how those reoccur. So in P2 "cucumbers are fruit" is again in green to show that it's the same statement as occurs in P1, and "tomatoes are fruit" in red: this is the consequent of P1. So our basic argument form goes P1, P2, we have an inference bar, and C for our conclusion. This word "if" signifies where the antecedent of P1 occurs, so let's write that in: P1. Then we have the word "then", we can put that in as well. This just shows that P1 is a conditional claim: the statement will have two parts. This "cucumbers are fruit" which I've underlined in green: we're going to use a letter for that. It's usual to start with P, so we'll put P, and we'll put it in green. Then we need to ask ourselves where else does that statement "cucumbers are fruit" occur in the argument, and it occurs in P2. And that is all of P2. So we just put the letter P next to P2 like that. Then we go looking for other propositions. "Tomatoes are fruit" is the consequent of P1, so we write Q to stand for that statement "tomatoes are fruit". We notice that that recurs and the conclusion, so we write Q for our conclusion as well. Now you can see that

the basic argument form here is "if P then Q, P, therefore Q" and that is our Modus ponens argument form. So because our original argument about cucumbers and fruit is of the Modus ponens form, we know that it must be valid.

Order of premises

You might wonder whether it matters what order the premises are in when you're identifying an argument form. And the answer is it doesn't matter at all. So let's look at a quick example of this. Here's our argument: "All events in the universe are causally determined. If all events in the universe are causally determined then free will is illusory. Therefore free will is illusory." We've got two premises and a conclusion. Our first premise is "all events in the universe are causally determined". Our second premise is the conditional one: "If all events in the universe are causally determined then freewill is illusory". And the conclusion is "therefore free will is illusory". So if we were to write down the form of that argument, the first premise is P on its own, the second premise is if P then Q, and the conclusion is going to be Q. So this is still Modus ponens, it's just put the premises in a different order from usual. But we can definitely identify it as Modus ponens. And you can see that this is going to work because if we were to swap the order of P1 and P2 we would have the same argument, -the premises would just be in a different order. And then that would put P in the second premise position in a way that we identify as being Modus ponens. So it doesn't make any difference what order the premises are in.

Chapter 3

Connecting premises

The purpose of this video is to give you some help on how to find the right sort of connecting premise to make an argument valid. Here are some general hints. An implicit premise or a connecting premise doesn't add any new reasons to the argument: it's just providing a connection between what's already there. If you find yourself adding new information to the argument in your connecting premise, then you're doing too much. Anything which is in the conclusion must be in the premises somewhere. So if something's mentioned in the conclusion, and it doesn't occur in the premises which are already present, then it needs to go in your connecting premise. Check that each thing in the premises connects to something else: you shouldn't have anything that's mentioned only once, because that means it doesn't connect to anything, and the idea is to connect together the parts of the argument. Okay, let's look at how some of these hints work in practice. So whatever is in the conclusion must be in the premises somewhere. Let's take this argument: "All politicians earn high salaries. Winston earns a high salary". So, we know that we have to get from premise one to the conclusion, and we're looking for what's going to need to go in P2 as the connecting premise. So we notice that "earns a high salary" is in the conclusion, but "earning a high salary" is also in premise one, so that doesn't help us at this point. However, "Winston"

is in the conclusion and "Winston" doesn't occur in premise one. So that tells us that we're going to need a premise that talks about Winston. Premise two is going to be about Winston. Then we look for what else has not yet been mentioned. "Politicians" is in premise one, but at the moment it doesn't occur anywhere else in the argument. That means that premise two is going to have to be about politicians so that politicians is connected in to the rest of what's said. Once we realize that we can see that the connecting premise needs to say "Winston is a politician". This argument is now valid: "All politicians earn high salaries. Winston is a politician. Therefore Winston earns a high salary". Let's look at another simple example. We need to make sure all of the parts of an argument connect. So consider this argument: "Daisy is a cow. Therefore Daisy is a mammal". Let's start by looking at the conclusion. We notice that the conclusion is about Daisy. "Daisy" is also in premise one: premise one is about Daisy. So probably the connecting premise doesn't need to be about Daisy. The conclusion says that Daisy is a mammal and the first premise says that Daisy is a cow, so we know that the connecting premise, premise two, needs to connect being a cow to being a mammal. Well it's clearly not going to say "all mammals are cows" because that would be false, and it wouldn't be useful – what the connecting premise needs to say is "all cows are mammals". So this argument is now valid: "Daisy is a cow. All cows are mammals. Therefore Daisy is a mammal". These examples are both very simple: they're just to get you started.

Let's look at a more difficult example. So here's our argument: "Premise 1: Making University education free will encourage lazy people to attend university. Premise 2: Lazy people will get no benefit from attending university. Conclusion: The government should not make university education free". So once again let's start with what the conclusion says. The conclusion is about making university education free, and making university education free also occurs in premise 1. So a connecting premise is probably not going to need to mention that. The conclusion is also about what the government should not do. Now neither the government, nor what should or should not happen, occur anywhere else in the argument. So the connecting premise is going to have to talk about what the government should not do. So let's write this in: "The government should not ..." Premise three is going to be about that. All right, and then we look at what else needs to be connected into this argument. So in premise one we talk about lazy people attending university, and premise two also talks about lazy people and their attending University. So those concepts are already connected together in the argument. Premise two also talks about getting no benefit from something, and that hasn't occurred anywhere else yet. So premise three is going to have to say something about not getting benefit. And in premise one what's still hanging out there, not circled, is this encouragement. So we have to connect what the government should not do to getting no benefits and encouragement, and we might have to think quite

carefully about how we're going to do that. But I know I'm going to want to say "the government should not encourage ..." And then I have to think quite hard about how to get the rest of the premise to say what I want it to say in order to make the argument valid. I want to say "the government should not encourage..." and then I have to think a bit about how to get the benefit in. It turns out I do end up repeating some phrases, because to make this argument valid I'm going to say "the government should not encourage people to attend university if they will not benefit from it". Now my argument is valid. Premise one says "making university education free will encourage lazy people to attend university". Premise two says "lazy people will get no benefit from attending university". And Premise three says "the government should not encourage people to attend university if they will not benefit from it". Then the conclusion follows "the government should not make university education free". This was a more complicated example, and sometimes you have to think quite hard about what that connecting premise needs to say. But with practice you'll get better at it.

Finding a more general connecting premise

When you're not sure how to add a connecting premise to make an argument valid, using a corresponding conditional is a guaranteed way to succeed. But it's often unsatisfying, and corresponding conditionals can be difficult to assess. In this example we're going to take you through how to use the technique of starting with a corresponding conditional in

order to find other, more meaningful premises to use as the connecting premise. So here's our argument. "Rents are rising to unaffordable levels. Housing is a basic need. The government should intervene." The first question to ask ourselves is 'what is the conclusion?'. Here there is no conclusion indicator, but as is quite common, the conclusion is the prescriptive claim: it tells us what should happen. So this is the conclusion. We add it at the bottom of the page: "the government should intervene". And we can see immediately that it's incomplete: the government should intervene in what? It's clear from the context, but when we put the argument into standard form we need to make it explicit: the way in which the government should intervene. So here we need something like "the government should intervene to stabilize rents at affordable levels". Then because we have our conclusion, we put in our inference bar, and then we look for our premises. The premises here are quite straight forward. "Rents are rising to unaffordable levels" is the first premise. And the second premise is "housing as a basic need". Now if we were going to use a corresponding conditional here, because we have two premises and then a conclusion, our corresponding conditional will just join both premises together with an "and", and there'll be an "if" before the premises, and the conclusion will go after the "then". So the corresponding conditional for this argument is "If rents are rising to unaffordable levels and housing as a basic need then the government should intervene to stabilize rents at affordable

levels". As I said corresponding conditionals are sometimes a bit satisfying. But we can use it. We know that this argument must be valid – it must be valid because it uses a corresponding conditional, and corresponding conditionals make any argument valid. So we can use this as a base to work out how else we might express the same points, and then we'll know that the resulting argument is valid. Okay, so we have this claim that says if rents are rising to unaffordable levels and housing is a basic need. So what's really important here is that there's a basic need and that basic need is becoming unaffordable. So the first part of the connecting premise just needs to say, "when a basic need becomes unaffordable..." Our consequent here says "then the government should intervene to stabilize rents at affordable levels". One of the things to ask yourself is, "is this specific to this government, or does it mean any government?" And I suspect that the person who asserts this premise means it not to be specific to our government but to apply to any government. So we just change "the government" to "a government". So "a government should intervene to stabilize rents at affordable levels", but we don't want to make this specific to rents anymore, this is the basic need that's become affordable. So, "when the basic need becomes affordable a government should intervene to stabilize that thing at affordable levels". So we know that this connecting premise must make the argument valid because it has just formed straight from the corresponding conditional. But it gives us a more general premise that somebody is more likely

to be happy to accept. So our connecting premise to make this argument valid is now "when a basic need becomes unaffordable a government should intervene to stabilize it at affordable levels". And this argument is valid.

Argument diagrams 1

This video introduces you to the basics of argument diagrams. It's often useful to diagram the structure of an argument, especially a complex argument. These are called 'argument diagrams' or 'argument trees' or 'tree diagrams'. I will take you through the basics of how this works. So when you want to show the structure of an argument, conclusions are always represented below the premises which support them, with a downwards arrow representing the inference: like this. Where the premises work together to support the conclusion they're braced together before the arrow points down: so you have premise one and premise two and you want to get to the conclusion which has always put below them. As premise one and premise two are working together, then we indicate that by joining them together before we have the arrow. Where premises or arguments threads support the conclusion independently, that's indicated by each having their own arrow to the conclusion. So suppose once again we have two premises and a conclusion, but in this argument premise one supports the conclusion on its own, and premise two supports the conclusion on its own. So our argument diagram looks slightly different. Let's look at some examples of how this works, and practice. Consider this simple argument:

premise 1 Daisy is a cow, premise 2 all cows are mammals, conclusion Daisy is a mammal. In this argument premise one and premise 2 are working together to yield the conclusion. Now when you have an argument that is this simple it's about hard to see why anybody would want to bother to diagram arguments. But as we look at a couple of more complicated examples, you'll be able to see why it's useful. So here's a longer argument. Every action I perform is either caused or uncaused. If my action is caused then it does not come about through an exercise of my free will. If my action is uncaused then it does not come about through an exercise of my free will. Therefore none of my actions come about through an exercise of my free will. If my actions don't come about through an exercise of my free will then I'm not responsible for them. Therefore I'm not responsible for my actions. So this is an extended argument, and we'll be able to map how the parts of the argument are related by putting it into a diagram. We can start off by noting that premises 1, 2, & 3 are all working together to get us to the first conclusion. So, we show that they're all working together – none of them are supporting that conclusion independently – by bracing them together with an arrow to get to C1. We can then see that C1 is operating also as a premise, and the premise that it's working with is P4. So C1 and P4 are working together to get to C2. So this is the argument diagram for this argument. The first three premises work together to support C1, and C1 works with P4 to support C2. It's a very common kind of diagram. Let's look at another example. If the reason

eating meat is wrong is because it curtails the life of a sentient being, then eating meat is not wrong so long as it doesn't curtail the life of the sentient being. Eating roadkill does not curtail the life of a sentient being. Therefore eating roadkill is not wrong. Eating animals which have died of natural causes does not curtail the life of the sentient being. Therefore eating animals which have died of natural causes is not wrong. Therefore not all eating of meat is wrong. Now the reason that argument diagram is going to help with this argument is that we really only have one way in standard form of dealing with complex arguments, and that's to stack all the premises one after another. But when we do a diagram we're able to show what the structure of this argument is. So P1 and P2 are working together to give us C1 in this argument. So that part of the argument structure is unsurprising. But C1 says eating roadkill is not wrong. We can see that the final conclusion, C3, not all eating of meat is wrong, follows from C1 alone. So C1 supports C3 on its own. So what's happening with the rest of the argument? P3 says eating animals which have died of natural causes does not curtail the life of a sentient being. This is not connected to P1 and P2 and C1. It is independent from that thread of the argument. So P3 does give a reason for us to accept C2. But P3 is not enough on its own for us to accept C2, so we know that P3 must be working with something else. So what goes over here? What is P3 working with? It's actually working with P1. If the reason eating meat is wrong is because it curtails the life of a sentient being, then eating meat is not

wrong so long as it doesn't curtail the life of the sentient being. And then P3 says eating animals which have died of natural causes does not curtail the life of a sentient being. Those two premises together support C2, eating animals which have died of natural causes is not wrong. I could have tried to diagram this differently so that P1 was attached by lines in two places, but in this particular diagram is just going to be cleaner if I write P1 down twice. And it is being used twice: It's being used in two different arguments threads. The final part of diagramming this argument is just to make C2 connect to C3, and once again C2 is a reason to accept C3 on its own. So now that we have this diagram completely mapped out, we can see something very interesting about the way that it works. Although both argument threads use P1, these are two separate arguments for C3: we have two independent reasons for why it's the case that not all eating of meat is wrong. This isn't really brought out by the standard form, and so sometimes an argument diagram is a very useful way of seeing what the structure of an argument is.

Argument diagrams 2

This video takes you through the way that you can use diagrams in a two-step process to help you put together your argument reconstruction. With a complex argument it's often useful to sketch out the structure of the argument before you begin your reconstruction. Then once you know the basic structure it's easier to put the argument into standard form, and to add any implicit premises. When you finish your

reconstruction you can do a final argument diagram which maps out the structure of the argument in a more complete way. Let's look at an example of this. Here's our argument: "Auckland will continue to have a housing shortage unless the construction industry can suddenly produce 20,000 houses. But that's not going to happen. It's too hard to get a building consent for starters, and they just don't have the workers. So the housing shortage isn't going anywhere." To help us deal with this argument I'm going to label each of the statements that are made in the argument. So our first statement is "Auckland will continue to have a housing shortage unless the construction industry can suddenly produce 20,000 houses". The next sentence is "but that's not going to happen". Then we have "it's too hard to get a building consent". And the next relevant sentence is "they just don't have the workers", so we'll call that '4'. Then '5' is "so the housing shortage isn't going anywhere". Once we've labeled all the parts of the argument we can look for the conclusion. The conclusion is very easy to find in this argument: it's got a nice conclusion indicator, "so", so we know that sentence five is going to be our conclusion. So that's going to go at the bottom of the page. And because it's a conclusion we know that it's going to be supported by something, so we mark that with an arrow. Now we need to think about the structure of the argument. So what is it that supports 5? 5 says that the housing shortage isn't going to stop. What is the main reasons that are given for why the housing shortage isn't going to stop? Here we can see that 5 is

supported by 1 and by 2. So 1 says "Auckland will continue to have a housing shortage unless the construction industry can suddenly produce 20,000 houses", 2 says "that's not going to happen". So what 2 really says is "it's not the case that the construction industry can suddenly produce 20,000 houses". So 1 and 2 are working together to support 5. Right. Then we will need to know what to do with the rest of the argument. Well, the two claims here are to do with why it's not going to happen, why the construction industry cannot suddenly produce 20,000 houses. So 3 "it's too hard to get a building consent", and 4 "the workers just aren't available", these are both reasons for why it's not going to happen. These are reasons for 2, and at the moment they're independent reasons: so there are two reasons why it's not the case of the construction industry can suddenly produce 20,000 houses. The first reason is "it's too hard to get a building consent" and the second reason is that "the workers just aren't available". So we've now mapped out the basic structure of the argument. This is going to help me a great deal when I put the argument into standard form, because I know that I'm going to want to start with sentences 3 and 4, and then I'm going to have sentence 2. Then I'm going to have 1 in order to get to 5. So I now know how to lay out my standard form: this has been made much easier because of my initial sketch. All right, so then we're going to speed up the reconstruction process. Here's what my standard form looks like. P1 It's difficult to get building consents in Auckland. P2 If it's difficult to get

building consents in Auckland then the Auckland construction industry cannot quickly produce 20,000 houses. P3 There is a shortage of construction workers in Auckland. P4 If there's a shortage of construction workers in Auckland then the Auckland construction industry cannot quickly produce 20,000 houses. Those reasons support C1 The construction industry in Auckland cannot quickly produce 20,000 houses. P5 Auckland will continue to have a shortage of houses unless the construction industry can quickly produce 20,000 houses. And then our final conclusion: Auckland will continue to have a shortage of houses. So now that we have the standard form laid out we can do a final argument diagram to show how the parts of the argument are structured. So we've got these four premises P1 – P4, but we can see quite clearly that P1 and P2 work together, and P3 and P4 work together. But these are separate argument threads. So P1 and P2 support C1 on their own, and P3 and P4 support C1 on their own: these are independent argument threads for the claim that the construction industry in Auckland cannot quickly produce 20,000 houses. Then C1 works with P5 to yield the final conclusion that Auckland will continue to have a shortage of houses. So here is our final standard form argument, and a final argument diagram. But the construction of the standard form that we have here was made much easier by the initial sketch that we did, which showed us which reasons we should put at the beginning of our standard form reconstruction. Here's a couple of hints for when you go to

do your own argument diagrams. When you do your own diagrams, if you have an arrow pointing anywhere but down, you've done something wrong. Arrows only point down because they always mark an inference to a conclusion. There's no need to do fancy computer graphics when you do your own argument diagrams for an assignment. It's often easiest just to draw the diagram quickly piece of paper, photograph it with your phone, and paste it into a document. When you do that your diagram will look something like this. It doesn't have to be flashy, it just needs to be clear how the structure of the argument works.

VERSIONING HISTORY

This page provides a record of changes made to this textbook. Each set of edits is acknowledged with a 0.01 increase in the version number. The exported files for this toolkit reflect the most recent version.

If you find an error, please contact sgibbons@waikato.ac.nz.

Version	Date	Change	Details
1.01	27 July 2024	Published chapters 1-6	
1.02			