

Lecture 3: Deduction and Induction

Lecturer: Right then, I think we ought to start. So welcome to everyone here and welcome to everyone on the podcast. Anyone who didn't get the thing I sent out last week about the argument that we tried to do, can you email me on that address and I'll bounce it on to you? I'll make sure that the thing that I sent out to everybody here is also available online for the people looking at the podcast later on.

Right, well, this week we're going to do different types of argument. (Slide 2) Whereas last week we learnt how to set out arguments logic book style. Do you remember, how was that done? What is it for an argument to be set out logic book style? Come on, I've got to get you warmed up, haven't I?

Female: Premise, premise, conclusion.

Lecturer: Premise, premise, conclusion. That's right. So you identify the premise and the second premise if there is one and the third premise if there is one, and then the conclusion. That's right, and the way I did it was to label them: 'Premise 1', 'Premise 2', 'Conclusion', but you could just put the premise 1, premise 2 and then a line and then the conclusion, there are different ways of doing it.

We dealt with ambiguities – hold on one second. How much did you miss? (Someone asked for the loop) ...

We dealt with ambiguities. Do you remember there were different sorts of ambiguity? What is an *ambiguity*, can anyone tell me?

Start off with what is an ambiguity?

Female: A [premise] with two or more meanings.

Lecturer: It's a word or a phrase that has two or more meanings, that's right. In the context where it is, it could mean two different things, which is very confusing. There are different types of ambiguities. Can anyone remember a couple of these types? No?

Male: Linguistic?

Lecturer: What do you mean by linguistic?

Male: Different meanings, for example: is.

Lecturer: Right, that's *lexical*. A lexical ambiguity is one word which could mean different things. So 'rum' could mean the drink or that it's strange or something like that. Or we use the example of 'bank' which could be all sorts of things: it could be a verb as well as a noun; all sorts of different things.

There are also structural ambiguities where you have: 'every pretty girl loves a sailor'. Does it mean '*there is a sailor such that every pretty girl loves him*' or '*every pretty girl loves some*

sailor', so there are different sailors who are loved by pretty girls, or just the one?

Then there's ambiguity of cross-reference. Does anyone remember what that is? Or *anaphora* you might have been taught in school? An anaphoric reference is when you have a pronoun that links back to a name. So: '*Jazz doesn't want James to come to the party because she doesn't like her*'.

Do you remember that? 'She' becomes ambiguous and the 'her' becomes ambiguous. Then there were ambiguities of *sound*, so the picture of water, do you remember, could be a pitcher or a picture? With a 't', oh, they've both got 't's, haven't they? A *pitcher* of water or a *picture* (drawing) of water?

We then learnt to identify conclusions and premises and we learnt there that although sometimes there are conclusion and premise indicators, what's the *only* fool proof way of identifying the conclusion of an argument? Can anyone tell me?

Male: The role in the sentence.

Lecturer: The role that it plays within the *argument* rather than within a sentence because, of course, the conclusion will *be* a sentence. That's exactly right. What role is played by the conclusion? It's the sentence that is...?

Female: Asserted.

Lecturer: That's being asserted, exactly so, and the premises which you can also tell in a foolproof way only by the role they're playing, are what?

Male: The reasons.

Lecturer: The reasons being given for the conclusion, exactly so. Then we went to eliminate irrelevancies, and that's where we got into a bit of hot water, isn't it, because you were all accusing me of throwing things out just for fun.

I hope that those of you who got the thing that I sent around see that what makes something relevant or irrelevant is the conclusion, because once you've identified the conclusion, that - you have the argument. Because the premises are the reasons for *that* conclusion, and anything that isn't a reason for that conclusion is irrelevant. Isn't it?

So even if it seems very important, like God, what could be more important than God perhaps? I'm afraid he gets left out if whatever is said about God is not a reason for believing that conclusion.

We identified suppressed premises and we saw that sometimes you have to put a suppressed premise in and sometimes you don't. When do you have to put it in?

Female: If it's controversial.

Lecturer: When it's controversial, good, that's right, because you can't leave it out of an argument if it's controversial. It would weaken

the argument to have it in, because it's controversial. When *can* you leave a premise out?

Female: When it's understood from the general context and it's obvious.

Lecturer: When it's blindingly obvious, when all you'd be doing is sort of eking out the arguments. I'm going to take my umbrella, do I have to say 'because it's raining and if I don't take an umbrella I'll get wet and I don't want to get wet'? No, of course I don't. You understand immediately when I say 'I'm going to take my umbrella'.

Finally, we learnt how to make terms consistent; what did that involve?

Male: Common language.

Female: [Depends on if it's raining].

Lecturer: That's right, if we had two words that meant the same thing or two phrases that meant the same thing, so 'it's raining' and 'it's pouring' we could just choose one of them and get rid of the other because all it does is confuse us.

We did the same thing with 'the cat is innocent' and 'the cat hasn't sinned.' I said that, in the context of this argument, there's no difference between them.

That's what we did last week: a bit of revision for us. (Slide 3)
So now we know what arguments are and we can distinguish arguments from other sets of sentences. Do you remember we can distinguish them from conditional sentences for example, and we can analyse arguments and set them out logic book style?

You're not doing badly, are you, for lecture three? (Slide 4) But we haven't yet started on how to *evaluate* arguments. That's important because some of you last week were having trouble with the analysis of an argument because as you were analysing it you were trying to evaluate it *as well*.

You were saying, 'This premise isn't true.' Well, maybe it's not but all we're interested in at the moment is whether it's a premise or not, we're not interested in whether it's true or not at the moment, but we *will* be.

So we're not even going to be able to start on analysis this week. What we're going to start on this week is looking at two different types of argument. It's very important to be able to identify the two different types of argument.

(Slide 5) We're going to look first at the *normativity* of critical reasoning, and of course I'll explain what 'normativity' means. We're going to look at the relation of '*following from*'. We're going to look at the two main types of argument and we're going to learn how to distinguish these two types from each other: deductive arguments from inductive.

(Slide 6) In critical reasoning we're concerned with whether or not an argument is good, so, importantly, it isn't a purely *descriptive* discipline. A purely descriptive discipline is one that tries to say what the world is like. So zoology for example, will attempt to describe different species, so it will attempt to categorise living things into species and into genus and so on,

and then it'll try and describe each species and so on, but it doesn't really get into '*shoulds*' in biology.

You might get into a should where all you mean by 'should' is you're talking about *normal*, something's being normal, so you're really talking statistics.

'Cats should have four legs' you might say. Well, yes, indeed, they should; statistically most cats *do* have four legs but this isn't a moral 'should'.

Whereas with critical reasoning we've got what's called a rational 'should'; it's *normative*, it lays down standards for us to follow. A good argument is an argument we want and we want it *because* it's good.

Critical reasoning is a normative discipline. We're interested in when a conclusion follows from a set of premises, and when it follows from a set of premises, it's good. So we're interested in arguments that are good, and that's when an argument follows from a set of premises, and we're also interested in which arguments are bad.

You need to be able to identify bad arguments as importantly as doing good ones. So following from is a very important relation between premises and conclusion.

(Slide 8) There are broadly two varieties of 'following from' and that's what we're going to learn to distinguish from each other today.

Both deduction and induction are varieties of 'following from'. So a conclusion may follow *deductively* from a set of premises or it may follow *inductively* from a set of premises. And really, these two types of arguments exhaust the types of arguments there are. There are arguments within each type but an argument is either deductive or it's inductive.

There are three questions you can ask to determine whether an argument is deductive or inductive. (Slide 10) Here are the three questions that you need to ask, and of course I'm going to look at each and explain them.

Anyone have any questions so far? No?

(Slide 11) The first property that distinguishes deduction from induction is the fact that a good deductive argument is truth-preserving, and bad deductive arguments aren't truth-preserving; inductive arguments are not truth-preserving at all.

So if an argument preserves the truth, it is a good deductive argument; if it doesn't, it may still be a deductive argument but it's a bad one; no inductive argument *ever* preserves the truth.

(Slide 12) What do I mean by an argument as being truth-preserving? An argument is truth-preserving if and only if, and that 'iff' is not me misspelling 'if' by the way. Whenever you see 'iff' – I suppose it could be an error – but it usually means 'if and only if.'

So you get the 'if' going in both directions, 'if P then Q', and 'if Q then P'. So an argument is truth-preserving *if and only* if it's not logically possible for its premises to be true while its conclusion is false.

So it's a very strong relation. This is the relation of *entailment*. So premises entail the conclusions when it's not logically possible for a set of premises to be true and a conclusion false.

(Slide 13) Before we look more closely at that, you've got to be able to distinguish *logical* impossibility from *physical* impossibility. This is actually crucial for a lot of philosophy, but it's particularly crucial for critical reasoning.

Female: Slower.

Lecturer: Am I going too fast? Okay. Let me go back to this one. (Slide 12) An argument is truth-preserving if and only if it's not logically possible ... (Slide 13) so this is why you've got to know the difference between logical impossibility and physical impossibility - if it's not logically possible for its premises to be true and its conclusion false.

So we've got to be able to distinguish logical impossibility and physical impossibility, and by the same token, logical and physical *possibility*.

(Slide 14) Something is *physically* impossible if it's inconsistent with the laws of *nature*, and it's *logically* impossible if it's inconsistent with the laws of *logic*. Something's inconsistent – two things (sentences) are inconsistent – if they can't both be true together.

The laws of nature are as they are and something is physically impossible if it can't be the case consistently with the way the laws of nature are. With logical impossibility it's the laws of logic that matter.

(Slide 15) So it's *physically* impossible for a human being to swim three miles underwater without breathing apparatus. Do you see that? If somebody told you they'd done that, you would know that they were not speaking the truth because it simply can't be done.

It's also physically impossible for a cat to talk intelligently about Kant. Even Oedipus, I've talked to him about Kant for years, he still doesn't talk back intelligently about it, or even unintelligently come to that.

But both these things are *logically* possible. Logic has nothing to say about whether my cat discusses Kant, and it has nothing to say about whether you can swim underwater for three miles without breathing apparatus.

(Slide 16) It's *logically* impossible on the other hand for a square to be a circle. So all those people, like my father, who talk about squaring the circle, it's nonsense; you cannot square a circle. Because if something's a square, it cannot be a circle, and you can tell that just from the very ideas involved.

A man can't be a married bachelor because if he's a bachelor he's not married, and if he's married he's not a bachelor. Just like if something's a square it's not a circle, and if it's a circle it's not a square.

We know that, and we know that a person can't have exactly three and exactly four children. These things are ruled out by *logic*. In other words, it's in the very nature of the concepts [square] and [circle], [married] and [bachelor], and so on, that we don't need to leave our armchair to know that there are no square circles.

This is why it's so comfortable to be a philosopher. (Slide 17) Here's an exercise. Give me two examples of states of affairs that are *logically* possible but not *physically* possible. Can you do that? Can you give me two examples of states of affairs that are logically possible but not physically possible?

If you remember, one example was that my cat can't talk intelligently about Kant, and another one is that a human being can't swim for three miles underwater without breathing apparatus. Can you think of another example like that?

Male:

The sun's so hot I froze to death?

Lecturer: The sun's so hot you froze to death? I'm not sure that isn't a logical impossibility because you're working on the words 'freeze' and 'hot', aren't you? They logically exclude each other, don't they?

Go on, have a go.

Female: 'Pigs can fly'?

Lecturer: 'Pigs can fly'. That's a very good one. Flying pigs, as far as logic is concerned, flying pigs are fine, but yes, there are no pigs that can fly. You might have been looking at the screen ... Any other examples?

Female: No, I didn't see that.

Lecturer: I believe that you believe you didn't see it but it's like you sort of see it and ...

Female: Yes, okay.

Lecturer: Can anyone think of another example of something that's *physically* impossible but *logically* possible?

Male: I can jump to the moon.

Lecturer: Good, it's logically possible for you to jump to the moon but it's *physically* impossible, isn't it? Yes, good. I think you're getting the idea. Anyone else want to have a go?

Male: I saw through the wall.

Lecturer: You saw through the wall. Good, yes, that's physically impossible but logically it's fine. There are lots of things like this that are physically impossible, ruled out by the laws of nature, but not ruled out by the laws of *logic*.

Of course there may be worlds out there that have different laws of nature, so when we want to know what's going on in those worlds it's to logic that we look, not to the laws of nature because they're not going to tell us what these worlds are like.

Can anyone give two examples of states of affairs that are *logically* impossible? 'You might have been frozen by a hot sun' might be a logical impossibility. Anyone think of another logical impossibility?

Female: It's raining cats and dogs.

Lecturer: Funnily enough, that's a *physical* impossibility; it's not a logical impossibility. Actually, it does rain frogs occasionally, I understand, it actually does.

Female: Burnt by ice?

Lecturer: It depends what you mean by 'burnt' there, doesn't it? But actually if you put your tongue on the ice in the Antarctic, you are in trouble, and it's actually called a 'burn', isn't it?

Female: A female boy.

Lecturer: Yes, that's a good example of a logical impossibility, yes. Anything else? It's quite difficult just thinking of examples but I hope you're getting the idea of the difference between physical and logical impossibility. Logical impossibility is inconsistent with the laws of *logic*: it's actually ruled out by the very concepts involved; and physical impossibility is ruled out by the *laws of nature*.

Whereas logical impossibility can be determined *a priori*, in other words we don't need experience of the world to determine what a logical possibility is, all you need is the concepts; with physical impossibility you need to know what the world is like.

Science tells us what's physically possible and impossible. What about these? By means of genetic manipulation we can produce pigs that are able to fly. Is that physically impossible, logically impossible, or are we unable to tell?

Female: Physically impossible.

Lecturer: It's physically impossible, that's right. Maybe by genetic manipulation, we will produce pigs that can fly one day, if we think we want them.

'John has exactly twice as many siblings as Janet; he has Susan and the twins'. Do you see how that's logically impossible? If somebody told you that, you'd think for a minute but then you'd think 'No, that's can't be right', and the reason it can't be right is if he's got *three* siblings, then he cannot have *exactly twice* as many as Janet has, however many Janet has.

'Muon neutrinos can travel faster than the speed of light in a vacuum'. Physically impossible, logically impossible?

Group: Physically impossible.

Lecturer: Yes, for a moment there we thought it might be logically possible and indeed *physically* possible, but somebody's microscope was shaking or something.

'Physicists have succeeded in building a time machine'. Is that physically impossible, logically impossible or such that we can't tell?

Male: Can't tell.

Lecturer: We can't tell on that one because it looks as if time travel is logically impossible, because if you could travel back and kill one of your parents before you were born, do you see the *logical* inconsistency that we get there?

If it's logically inconsistent, it shouldn't be *physically* possible, but there are physicists who argue that it's not only *not* logically impossible, it's *physically* possible. If you look on YouTube, if you put in 'time travel - physically possible', you'll get up a physicist who believes that it is physically possible. So I'm not going to say. I think I'm inclined to think it's logically impossible.

Male: Is there a point about all of this in the sense of you have to have, almost as a caveat or let's say, 'given what we know today'?

Lecturer: With physical impossibility that's exactly what we have to say, yes. Because the laws of nature, we may be able to work with them to produce a pig that can fly. Obviously we can't work against them, but as things currently stand, flying pigs are against nature.

But logical impossibility, no; logical impossibility seems more fixed because the laws of logic seem more fixed. But quantum mechanics, for example, has called into question a couple of the laws of classical logic.

So whereas there's a distinction between the laws of nature, i.e. the natural uniformities that govern our world and the laws of nature that we put together, i.e. our *descriptions* of the natural uniformities, do you see the difference? So there are the laws, our laws, if you like, and the natural uniformities that we describe.

In the same way the laws of logic, there's our *description* of these laws and there are the laws themselves, and we may have the wrong description and may need to tighten it up at

some point. So we can't be sure we've got the right understanding of either type of law.

But it's certainly true with physical impossibility that it's malleable according to our physical knowledge.

Let me just ask, are there any other questions about physical and logical impossibility before I move on? Because it's quite an important distinction.

Male: Would Quine have produced an argument against logical certainty? He certainly challenges it.

Lecturer: No. What he challenges is the existence of 'analytic truths'. So he doesn't think that there are truths that are not revisable

Male: Okay.

Lecturer: (slide 18) An argument is truth-preserving if and only if it's not logically possible for its premises to be true and its conclusion false. That's the definition we had up there before.

(Slide 19) It's important to notice that the actual truth value of the premises of the conclusion is irrelevant to determining whether an argument preserves truth. (Slide 20) Have a look at this argument: 'All heavenly bodies revolve around the Earth'. 'The Sun is a heavenly body.' 'Therefore the Sun revolves around the Earth'.

Do you see that if those premises are true, the conclusion *must* be true? Do you see that? But the conclusion is false, as is at least one of the premises. Do you see that?

So we have an argument that *preserves* truth but doesn't *generate* truth. That's the importance of the distinction: you've got to distinguish between an argument *preserving* truth and generating truth. If it preserves truth, what you're saying is 'if there is truth in the premises there will be truth in the conclusion'. That suits this argument very well.

If these premises were true, the conclusion *would be* true; it would *have to be* true, wouldn't it? It couldn't possibly be false. It's logically impossible, in fact, for that conclusion to be false if those premises are true, and that's what it is for an argument to be truth-preserving.

If an argument has that property, then it's called 'valid', as we'll discover next week, and the premises *entail* the conclusion. This is a good deductive argument, and that's so even though the conclusion is false. We'll look at this a bit next week but can anyone tell me why the property of being 'truth-preserving' is useful, even though you can end up with a false conclusion? Why should we *care* that our argument should be truth-preserving even though you could end up with an argument with a false conclusion?

Male: Because you can trace back to find the false premise in that.

Lecturer: Good. If an argument is truth-preserving, you know *either* that its conclusion is true if its premises are true, *or* you know that if the conclusion is false, at least one of the premises must be false, and of course, you can look back to find which premise is false.

Of course, if this is 'all swans are white,' the thing in the Swan River in Australia, if it's a swan it will be white, but it's *black*,

you know that one of your premises is false. So it can be empirically very useful to have the property of truth-preservingness in an argument.

(Slide 21) We can be certain of an argument that preserves the truth that it's a 'good deductive argument'. It's not just deductive, it's also a good deductive argument if it preserves the truth.

(Slide 22) We can be just as certain that if an argument *doesn't* preserve the truth, then it's either a *bad* deductive argument or it's an inductive argument. So it could be one or the other, but it's certainly not a good deductive argument.

(Slide 23) Let's do this exercise. Which of these following arguments is truth-preserving, in other words, is a good deductive argument? Have a look at them yourself and put your hand up when you think you've got the answer to one, let's say.

Hands up when you think you've got the answer to one.

'Tom is a banker', 'all bankers are rich', 'therefore Tom is rich'.

Lady there.

Female: That's a good deductive argument.

Lecturer: Does everyone agree? Yes? You're absolutely right. If Tom is a banker and if all bankers are rich, then Tom *must* be rich, mustn't he? There's no question about it. So that's a truth-preserving argument and it's a good deductive argument therefore.

What about: 'Sue and Tom lead similar lives' but 'Sue smokes and Tom doesn't'. 'Therefore, Sue is more likely to die from heart disease than Tom'?

Put up your hand when you think you know whether that's truth-preserving or not.

I'll let a few more put their hands up first.

Do you want to tell us, do you think it's truth-preserving?

Female: No, because I think it's missing something. As it is, no, but I'll stop there.

Lecturer: As it is, no, and that's because there is a possibility that the conclusion is false even if the premise is true. Why is that?

Female: Because it doesn't make any link between smoking and heart disease: it needs an extra premise.

Lecturer: Right. It could be a premise that we just want to leave out or it could be a controversial premise. Don't let's follow it. You're quite right, there is something left out and that premise could be true and that conclusion false, so we know that that argument doesn't preserve the truth, so there could be premises but no truth in the conclusion, so it's not truth-preserving. So either that's a bad deductive argument or it's an inductive argument, and we'll see later which it is.

What about, 'All dogs are mortal'. 'Lucy is mortal.' 'Therefore Lucy is a dog'? Put up your hand when you think you know. Gentleman there.

Male: Not.

Lecturer: It's not truth-preserving, that's right. It's a *bad* deductive argument this one, isn't it? We'll see why it is later on. But yes. 'All dogs are mortal'. 'Lucy is mortal.' 'Therefore Lucy is a dog'. Those two premises can be true and yet that conclusion false quite easily, couldn't they? Lucy might be human and mortal. What about No.4? Have a look at that and put your hands up when you think you know whether it's ...

Gentleman there.

Male: It's a good deductive argument.

Lecturer: Everyone else agree?

Male: No.4, yes.

Male: Sorry, No.4.

Lecturer: One, sorry. I thought you were wrong, but actually I think you're right.

'Killing is wrong.' 'Therapeutic cloning involves killing.'
'Therefore therapeutic cloning is wrong'.

Yes, exactly. That's a good deductive argument. That if it's true that killing is wrong and it's true that therapeutic cloning involves killing, then it must be true that therapeutic cloning is wrong.

If you want to deny that, you'd have to go back and deny one of the premises, wouldn't you? You'd have to say one of those premises is false.

What about No.5? Hand up when you think you've got ... I've given this one away, haven't I?

Female: Yes.

Male: It's a bad deductive argument.

Lecturer: No, all you need to say at the moment is it's not a good deductive argument, it's not truth-preserving. Actually, it's not a deductive argument, as we'll see later on.

'Every person with Huntington's disease who's been examined has had the HD gene on chromosome 4.' 'Therefore everyone with HD has the HD gene on chromosome 4.'

That (premise) could be true and that (conclusion) could be false. It's true that it's unlikely to be false but it *could* be false, couldn't it? It could be that we just haven't examined any of the HD people who've got the HD genome on chromosome 5 perhaps, or whatever. If there are any biologists in the room ...

What about this one? 'If the liquid is acidic, it will turn litmus paper blue.' 'This liquid does not turn litmus paper blue.'

'Therefore, this liquid is not acidic.' [Everyone likes to point out

that acid turns litmus paper *red*, but this is irrelevant to whether the argument is a good one or not! *MT*]

Hand up. Lady there.

Female: Me? It is a good deductive argument.

Lecturer: Everyone agree? Good. We've got a feel for when an argument is truth-preserving or not and we know that if it's truth-preserving, it is a good deductive argument, and if it's not truth-preserving it's either a bad deductive argument or it's an inductive argument.

That's the first question that we ask to distinguish deductive arguments from inductive. (Slide 24) But it would be reasonable to ask at this point perhaps: why is deduction useful? (Slide 25) If we believe the premises of a good deductive argument and it's logically *impossible* for the premises to be true and the conclusion false, then to believe the premises is to believe the conclusion, isn't it?

Do you see that in a good deductive argument, the conclusion doesn't go any further than the premises? What is contained in the premises is contained in the conclusion. So in a way, you learn *nothing* from a good deductive argument.

(Slide 26) It doesn't tell us anything we don't already know. It gives us the logical consequences of what we already know.

(Slide 27) If human beings were perfectly rational and believed all the logical consequences of all their beliefs, then deduction wouldn't be useful. The reason it is useful is because we're not perfectly rational. (Slide 28) We often need the consequences

of our beliefs spelled out to us, and that's what deduction does, and I just love this sentence, so I put it in.

I slightly changed it to make it even better but it came from Mark Sainsbury.

A deductive argument, if the premises are true, then the conclusion must be true, there is no logically possible situation in which the premises are true and the conclusion false.

That's a good deductive argument. One question...

Male: Sorry, just back a little bit. On the need for a conclusion or not, does the conclusion not link the two premises in a way that wouldn't have happened without the conclusion?

Lecturer: Yes. The conclusion is a logical consequence of the two premises taken together.

Male: No, more than that, it links them.

Lecturer: It links them because you've got to put them together in order to get the conclusion. If you have 'all heavenly bodies revolve around the Earth –'. Is that what I had? I can't remember what it was.

[Cross talk 0:36:00].

Male: That was one that wasn't logical.

Lecturer: No, it is logical, this is truth-preserving. (Slide 20) But if you only have that premise without that one, then that one wouldn't logically follow, would it?

Male: You were saying it didn't *advance knowledge* any more, but it linked premise one and premise two in a way that wouldn't otherwise have happened.

Lecturer: That's certainly true. The fact is that that belief is a logical consequence of that belief (premise one) and that belief (premise two). So it links them because you need both those beliefs to get that. Does that satisfy you?

Male: Yes.

Lecturer: Good. Where were we? (Slide 29) The second question that we ask to distinguish deductive arguments from inductive arguments is: 'is the matter of whether the conclusion follows from the premises an either/or matter, or is it a matter of degree?'

Does the conclusion follow from the premises either definitely or definitely not, or could you say 'well, maybe it does and maybe it doesn't'? Let's have a look.

(Slide 30) If it's an either/or matter, the argument is deductive. If it's either good or it's bad, then it's a deductive argument. If it is a matter of degree – if it's, 'No, it's not bad, it's better than the last one,' or something like that – then it's an inductive argument.

(Slide 31) We know that the conclusion of a deductive argument follows from its premises if and only if the argument is truth-preserving. So if a deductive argument is not truth-preserving, its conclusion doesn't follow from its premises. It's either it does or it doesn't. It's either truth-preserving or it's not. It can't be anything in between.

(Slide 32) The evaluation of a deductive argument is an either/or matter: either the conclusion follows from the premises or it doesn't; there's no maybe, no degrees of following from at all when we're talking about deduction.

(Slide 33) Here's a good argument: 'All bankers are rich.' 'Deepak is a banker.' 'Therefore Deepak is rich.'

Do you see that if those premises are true, the conclusion must be true?

Here's a bad one: 'All bankers are rich.' 'Deepak is rich.' 'Therefore Deepak is a banker.'

Do you see that those two premises could be true without the conclusion's being true? We can see that immediately, can't we? There's no 'maybe' here, one is good, one is bad and there's nothing in between, and that's always the case with deductive arguments. Question over here.

Female: On the killing one, killing is wrong ...

Lecturer: Therapeutic cloning?

Female: Yes. It didn't say 'or'.

Lecturer: Let me go back and see what it did. I think it implied 'or' but let me see.

Female: If you were arguing with someone on that one you'd say –

Lecturer: Killing is wrong, therapeutic cloning involves killing, therefore therapeutic cloning is wrong.

You mean this doesn't say '*all* killing is wrong'?

Female: Yes.

Lecturer: No, but it does imply it, doesn't it? It says quite generally 'killing is wrong.'

Female: Right, there's no degree there?

Lecturer: The only way you could make that conclusion false is by making that not true. In other words you'd have to say 'some killing is wrong, most killing is wrong, but it's not the case that *all* killing is wrong.'

Female: Okay.

Female: Marianne, can I just ...

Lecturer: That was a good objection.

Female: Just to follow up, if you then said there's a question of degree around the therapeutic cloning, sometimes it involves killing...

Lecturer: Then you would be falsifying that premise. With this argument, it's a truth-preserving argument in that *if* that premise is true, *and* that premise is true, then that conclusion *must* be true, and if you don't think that conclusion is true, you've got to question either that premise or that premise, and you could question them in the ways that you've just discussed.

Any questions about that? No? Let's go on to – now I've forgotten where we were. Ah, yes.

So you can see one is good, one is bad, and that's always the way it's going to be. (Slide 34) So when a deductive argument is good, when it's truth-preserving, i.e. when its conclusion follows *deductively* from its premises, its premises are *conclusive* reason for believing its conclusion. It's either good or its bad, so if it's good, then its premises are *conclusive* reason for believing the conclusion.

(Slide 35) That's because deductive arguments are 'monotonic'. Here's a new word for you, at least I assume it's a new word, maybe it's not. Deduction is monotonic, and what that means is that we can add anything we like to a good deductive argument without making it bad.

You'll also see that we can add anything we like to a *bad* deductive argument without making it good!

If a good *deductive* argument is good, it's conclusively good. [I said here that 'if it's bad it's conclusively bad'. But that isn't

true! M.T.] So there's nothing we can learn that would change a good deductive argument to being a bad one.

(Slide 36) Let's have a look, here's a good deductive argument. Would you like to have a go at giving me a premise that might change that from a good argument to a bad argument? Have a think about it.

Female: 'Some bankers are good bankers.'

Female: 'All *good* bankers are rich.'

Lecturer: That's not adding a premise, that's changing one of the premises that's there, which is not playing the game, sorry.

Female: 'Some bankers are rich.'

Lecturer: Okay, you're adding that... I'll tell you what, let me show you what I've got, what I thought you might come up with, because I think that might be easier.

(Slide 37) What have I tried to put in there? I can't remember now. [Cross talk 0:42:39]. I tried to add 'it's not the case that Deepak is a banker' here. So there's the argument, the first premise, second premise and conclusion, and I'm adding that in (third premise) in the hope that that might make it a bad argument. Do you think it does make it a bad argument?

Male: No.

Lecturer: It doesn't, does it? It certainly doesn't change it from 'if all these premises are true, this conclusion is true,' does it? Because, actually, if all these premises are true – anything follows from a contradiction – and we've got a contraction in here, haven't we? Do you see? So anything follows from it.

What about this one (right hand argument), what have I added here? 'No banker is rich.' So I've contradicted that one. But we've got exactly the same thing, haven't we? It hasn't gone from being the case that *if* all the premises are true, the conclusion must be true again.

So with deduction, you've got absolute - it's *conclusive*, it's either good or it's bad and there's nothing you can do.

Male: I'm not sure I understand that.

Lecturer: Let me do the bad one and see if that helps. If not, we'll go back.

[During the lectures I was having something of a brainstorm! It is quite true that deductive validity is monotonic. But *invalidity* isn't. You *can* change a bad deductive argument into a good one by adding further premises. Inductive arguments can be changed from good to bad *and* bad to good by adding further premises ... MT]

... With a good argument, nothing you could add would change it from being truth-preserving to being not truth-preserving.

Male: I'm okay.

Lecturer: You're okay now?

Male: Where I was confused is I thought that you were still saying it's a good argument and I wasn't seeing that, I think it was a bad argument.

Lecturer: What you're not doing is changing it from a good argument to a bad argument.

Male: Exactly.

Lecturer: Let's have a look at this deduction. This is what I'm claiming. If a deductive argument is good, its premises are *conclusive* reason to believe the conclusion. There's nothing we're going to learn that's going to change that.

When a deductive argument is bad, then its premises are *no reason whatsoever* to believe the conclusion.

That doesn't mean that the conclusion of a bad deductive argument might not be true. Do you see that? It means only that its premises are no reason to believe the conclusion.

(Slide 38) Here's a *bad* deductive argument: 'If it's Monday, Marianne will be wearing jeans.' 'Marianne is wearing jeans.' 'Therefore it's Monday.'

It is a bad deductive argument, isn't it?

Group: Yes.

Lecturer: Good, thank you. I've got to the point where I can't...
But is that conclusion false?

Group: No.

Lecturer: It isn't, is it? No, it is Monday. But our reason for believing it's Monday doesn't depend on that argument, does it? Or at least I would hope so, because if it depended on that argument ...

So you can have a bad deductive argument with a true conclusion in exactly the same way you could have a good deductive argument with a false conclusion.

If you have a good deductive argument with a false conclusion you'd know *what?* What do you know if the conclusion of a good deductive argument is false?

Group: One of the premises is false.

Lecturer: Good. Well done.

Male: Can I ask a question? If you put in an extra F after the first F, if you gave it one of your double F 'iffs', does that change it to a good deductive argument?

Lecturer: That would change it completely. No it would then say 'if and only if it's Monday Marianne will be wearing jeans.' 'Marianne's wearing jeans.' 'Therefore it's Monday.' I think that would make it a good argument.

Male: That's what I thought.

Female: But you're changing the premise.

Lecturer: But you'd be changing the premise.

Male: Yes, just curious to know that if you just changed one tiny bit of it ...

Lecturer: Yes, but making an 'if' into an 'if and only if' is not making a tiny change, it's making a huge change, logically.

So either a deductive argument is good because its premise is a conclusive reason to believe the conclusion: it's truth-preserving; or a deductive argument is bad because its premises have no reason at all to believe its conclusion.

With deduction, it's always an either/or matter. But no inductive argument preserves the truth, so we can't evaluate inductive arguments by appeal to whether or not they preserve the truth.

The very strongest inductive argument still doesn't preserve the truth. So we evaluate them according to how strong they are, and of course strength is a matter of degree, isn't it?

(Slide 42) Here an inductive argument is strong if and only if the truth of its premises make its conclusion *significantly more likely* to be true; and an inductive argument is weak if and only if the truth of its premises makes its conclusion only *slightly more likely to be true*. Here are some examples.

Who thinks the one on the right-hand side is the strong one?
No, that's this one.

Female: That one?

Lecturer: That one. Yes, you're right. Okay, who thinks the one on the left-hand side is strongest?

That's a very strong inductive argument, isn't it? 'The sun has risen every day in the history of the universe' suggests that it's highly likely it's going to rise again tomorrow, isn't it?

Whereas 'every time I have seen Marianne she has been wearing earrings.' 'The next time I see Marianne she'll be wearing earrings'. Okay, it's probably true isn't it, that every time you've seen me I've been wearing earrings, but you've only seen me three times, some of you. Is that really enough? Maybe next time you call around to my house at 4 o'clock in the morning, get me out of bed, I promise you I don't wear earrings in bed.

So that's a fairly weak inductive argument. But do you see they both have the same form, if you like? So in each case you're saying 'all the things that I've seen like this have been this, therefore the next one, the one I haven't seen will be like this.'

(Slide 43) So inductive strength is a matter of degree because whether the premises of an argument make the conclusion more or less likely is itself a matter of degree.

Inductive strength is *not* monotonic. I taught you that new word. Deduction is monotonic. What that means is you can add anything you like to a good argument, you can't change it from being good; you can't stop it being good.

(Slide 44) But inductive strength is not monotonic, the addition of a new premise to an inductively strong argument might make it inductively weak, and I've got some examples coming up.

Also, the addition of a new premise to an inductively *weak* argument might make it inductively *strong*, so let's have a look at how this works.

(Slide 45) Take a relatively strong inductive argument: 'Jones confessed, therefore Jones is guilty'. You agree probably, that's a pretty strong inductive argument. But then add: '10 independent witnesses testify to Jones being 100 miles away from the scene of the crime at the time it was committed'. Suddenly it looks very weak, doesn't it?

What we've done is we've taken a strong inductive argument, we've added another piece of information and we've weakened the argument completely, haven't we?

Group: Yes.

Male: Jones is a twin.

Lecturer: Do you remember when I said about ambiguity? We've got to assume that Jones is referring to the same person in the same context.

(Slide 46) Or take the relatively weak inductive argument: 'Jones was present at the scene of the crime therefore Jones is guilty of the crime'. Add to the argument which I think you all agree is weak inductive argument, yes?

Group: Yes.

Lecturer: Add the extra premise: 'Smith the policeman who tried to stop Jones kill the man, saw Jones plunge the dagger in his heart' (I was having fun with that one).

Suddenly a weak argument looks very much stronger, doesn't it?

Group: Yes.

Lecturer: So induction is not monotonic. You can strengthen and weaken inductive arguments by discovering new information. You can't do that with a deductively valid argument; with deduction, if an argument's good it's good and that's the end of it.

Male: Are you in danger there of turning an inductive argument into a deductive argument by the addition of the additional statement?

Lecturer: It's very easy to turn inductive arguments into deductive arguments by adding the right premises. Have I done that here?

Male: I would have thought you were in danger of doing that in the second statement.

Lecturer: Well, no, not if Jones died of electrocution, and he might have done, mightn't he, consistently with that? I hadn't actually said what the crime is here.

Male: Okay.

Lecturer: You can come back to me on that one if you want to later on.

(Slide 47) Let's have a look at this. Exercise three, which is very similar to exercise two. What we want to know now is whether it's being good or bad as a matter of *degree* or a matter of either/or.

Whereas last time you identified the *good* deductive arguments, how you can also identify the *bad* deductive arguments and distinguish them from inductive arguments:

'Tom is a banker.' 'All bankers are rich.' 'Therefore Tom is rich.'

What's that?

Group: Deductive argument.

Lecturer: Good, deductive, as it was before, yes, it doesn't change:

Sue and Tom lead similar lives but Sue smokes and Tom doesn't. Therefore Sue is more likely to die from heart disease than Tom.

Stick your hand up when you think you've got the answer to that.

What we want to know, we know that it's not a good deductive argument from the last time, but is it a *bad* deductive argument or it is an inductive argument?

Female: Inductive.

Lecturer: It's inductive, exactly so. What could we add to that to make the conclusion much less likely?

Female: Why would you want to make it less likely?

Lecturer: Just because if we can, we know it's inductive. So what could we add to that?

Female: You could add that Tom's got a bad heart.

Lecturer: Or that all of Tom's family died young of heart disease, yes, that would immediately change the strength of that argument, wouldn't it? As it is, it looks pretty strong, but we can weaken it immediately by giving a bit of Tom's family history. Is that a question?

Female: When it suggests that Sue and Tom lead similar lives, probably if you leave the first statement out, it might be stronger because Sue smokes and then Tom doesn't, so obviously Sue is more likely to die.

But when it suggests that Sue and Tom lead similar lives, it's not true, because she smokes and he doesn't. Isn't it?

Lecturer: They could lead similar lives apart from that. They could both exercise, both eat lots of fruit and vegetables, both not eat much animal fat.

What about, 'All dogs are mortal.' 'Lucy is mortal.' 'Therefore Lucy is a dog'? Put your hand up if you think you know whether that's a bad deductive argument or an inductive argument.

Male: A bad inductive argument.

Lecturer: It's bad inductive argument. How do you know that?

Male: Because of the premises.

Lecturer: Tell me more.

Male: It's not truth-preserving, is it?

Lecturer: It's not truth-preserving, and in fact, the premises are no reason whatsoever for believing the conclusion, are they? So whereas with the first one, the premises are conclusive reason for believing the conclusion, with the third one the premises are no reason whatsoever for believing the conclusion. That's a bad deductive argument.

What about number 4? Well, we know that's a good deductive argument, don't we? Okay, what about number 5? Is that a bad deductive argument or an inductive argument?

Group: It's inductive.

Lecturer: It's inductive, that's right:

'Every person with Huntington's disease who's been examined has had the HD gene on chromosome 4.' 'Therefore everyone with HD has the HD gene on chromosome 4.'

Could you add anything to that to make it a bad inductive argument? As it stands it looks pretty strong. What might we do to change it into a weak one?

Male: We've only examined five people.

Lecturer: Exactly. If we've only examined five people, then it suddenly looks a lot less strong. Good.

The last one was a good deductive one, wasn't it?

Group: Yes.

Lecturer: Good. (Slide 48) So we know now that if an argument preserves the truth, it's a good deductive argument. If an argument doesn't preserve the truth it's either a bad deductive argument or it's an inductive argument. We know that if an argument is either good or bad with no matter of degrees or no maybes, then it's a deductive argument.

If it's been good or bad as a matter of degree, then it's an inductive argument. We've looked at our first question and our second question, now let's look at the third.

(Slide 49) The third question was: *can we evaluate the argument a priori?* (Slide 50) To be able to evaluate something *a priori* is to be able to say whether it's good or bad, or true or false, or whatever, something normative, without experience, without needing to go and look at the world, if you like, you can do it simply from logic on its own.

The opposite is *a posteriori* which is when you have to actually go and look at the world. So let's have a look at this.

To be able to evaluate an argument *a priori* is to be able to tell whether the argument is good or bad by appeal *only* to the structure of the argument and to the logical words used in it without need of any information about the world.

You might want to say you need information about language and that indeed is true, so you need some information about

the world, but you don't need more than information about language and meanings.

(Slide 51) So these are the logical words: 'and', 'if then', 'not', 'or' and 'if and only if'. There are other logical words but we won't worry about those at the moment. (Slide 52) This gives you an indication of what's meant by the *structure* of an argument, so (Slide 52) here's an argument we used last week: If it snows, the mail will be late. It is snowing, therefore the mail will be late.

Can you see that if you were to take, 'It is snowing,' as 'P', then you've got 'if P' and you've got 'P' here again, and 'the mail will be late' becomes 'Q'. So you've got: 'if P then Q', 'P' therefore 'Q'.

Do you see the structure of the argument can be brought out just by taking away the sentences in the argument leaving the logical words.

(Slide 53) Is this a good argument? Put your hand up when you've decided, rather than yelling out.

[pause]

Put your hand up when you've decided, up so I can see it because I'll ask when there's a critical mass of people who know, or think they know.

What do you think: is it a good argument?

Female: Yes.

Lecturer: It is, isn't it?

Female: 'P then Q'.

Lecturer: Does anyone know what it means, does anyone know which it is or what it is for something to be '*pomol*'? No? Good, I'm not surprised because I made them up.

But you see the content here is completely irrelevant to your ability to evaluate this argument. You really don't need to know anything about the subject matter in order to be able to say that that argument is good.

The reason for that is that what makes that argument good is the way the logical words in it are combined, it's the *structure* of the argument, it's nothing to do with the *content* of the argument, and that's true of all deductive arguments. All deductive arguments can be evaluated a priori.

(Slide 54) Here's another one, is this a good argument?

Group: No.

Lecturer: It isn't, and you can tell that because this is actually a bad deductive argument, whereas the last one was a good deductive argument.

(Slide 55) It's interesting to note that the fact the deduction is a priori actually means that it's *topic neutral*. This is what makes logic, or deductive logic anyway, such a useful transferrable skill because it doesn't matter what you're talking about, what you're thinking about, if you can recognise a deductive argument, you will be able to evaluate it even if you know nothing whatsoever about the subject matter.

The subjective matter is completely irrelevant to evaluating a deductive argument because what matters is the logical words and the way they're combined.

Male: Excuse me. In the first premise –

Lecturer: Who's speaking?

Male: I am. In the first premise ...

Lecturer: ... Of the last one?

Male: The nonsense words.

Lecturer: That one?

Male: If 'widgets are pomol' meant positive and 'havena' meant negative, then the first premise in those terms stands up but it's nonsense.

Lecturer: What you're doing is you're trying to put content in that's going to make it nonsense, but it's already nonsense, isn't it? You don't need to do anything to make it nonsense, it's nonsense already, but it's still a good argument or a bad argument or whatever it is.

Male: But supposing it was code: 'If widgets are positive, then widgets are negative.' That is a premise that cannot stand up in logic.

Lecturer: No, it's the content. What you're saying is: 'it's false'.

Male: Yes.

Lecturer: Actually, it's not even false because it is meaningless, isn't it? 'If widgets are positive, then widgets are negative.' That's nonsensical, it doesn't have meaning.

Male: Yes.

Lecturer: But you only know that from the content that you're putting in there, and my aim in using this is to show that you can decide whether something's a *good argument* or not without knowing what the content is about.

So wanting to put in some content that doesn't make sense is not changing my argument at all.

Male: So you could use positive and negative if you didn't know what positive and negative meant? [Inaudible discussion]. Why would you, but you could.

Lecturer: What my aim is to show you is that you can tell whether a deductive argument is good or bad without knowing what the content means.

I take it that because you were able to see that the first argument was good and that the second argument was bad, I take it that I've proved my point?

Female: Yes.

Lecturer: I could go back and discuss – but I'm inclined not to – because I don't see why changing one nonsensical argument into another nonsensical argument is going to gain us anything.

May I move on?

Group: Yes.

Lecturer: I'm not sure I've dealt with you, perhaps we can talk afterwards.

Male: Marianne, if you replace the 'P's and 'Q's, then it's clear.

Lecturer: It's clear even without the 'P's and 'Q's.

Male: It is, but you're saying all 'Ps are Q', 'X is a Q' –

Lecturer: Let me carry on because I'm now going to replace things with 'P's and 'Q's, or at least I can show you how.

(Slide 56) Do you see that these two arguments have exactly the same structure? Anyone want to deny that? Can you help me work out what the structure is using 'P's and 'Q's? Let's use that one.

'If P then Q', 'P' 'Therefore Q'. So here it ('P') means 'it is snowing' and 'Q' means 'the mail is late'. What's 'P' in the second argument? [inaudible discussion], but what's 'Q'?

Group: It describes it.

Lecturer: It's the act is to remove the cross-reference, the act is right. Yes, exactly. So we've got exactly the same structure it does for both those arguments, doesn't it?

Female: Why do you use 'P' and 'Q'?

Female: It's tradition.

Lecturer: Instead of 'R' and 'S' or something, yes, it's just convention.

Female: Is it convention to do with pigs and queens?

Lecturer: No, not that I know of. Actually, some people would use 'S' and 'M' here just because 'snowing' begins with 'S' and 'mail'

begins with ... But it would be then less relevant to that one, so it's convention that we use 'P' and 'Q's.

(Slide 57) Hang on, here we've got another one:

'All widgets I've ever seen have been havena.' 'Therefore all widgets are havena.'

Can we tell whether that's a good argument or not? We can't, can we? Why not? [inaudible discussion].

You don't know how many widgets you've seen, you don't know what widgets are and you don't know what it is to be havena. You do need to know what the content is and you need to know something about the content in order to be able to evaluate that argument.

So with a deductive argument, it can be evaluated *a priori*, i.e. on the basis only of its structure in the logical words. If it can only be evaluated by appeal to some background information, some information that you have, then it's an inductive argument.

(Slide 58) Inductive arguments can be evaluated only *a posteriori* in the light of an understanding of the content of the argument and by bringing to bear background information about the world.

Whereas you can strip the content out of the deductive argument and still determine whether it's a good or bad argument, there's no way you can strip the content out of a good inductive argument and still hope to evaluate it.

(Slide 59) Can we evaluate these arguments *a priori* or not?

'Jennifer is tall.' 'Jennifer is the bank manager.' 'Therefore the bank manager is tall.'

Put your hands up when you've decided whether it can be evaluated a priori.

What do you think?

Female: Yes, it can.

Lecturer: Yes, it can. It's the structure of that argument that tells us whether it's a good argument or not. Is it good?

Group: Yes.

Lecturer: Yes, it's a good argument, isn't it? What about, 'Crocodiles are dangerous.' 'James's dog is dangerous.' 'Therefore James's pet is a crocodile'?

I think I was getting tired at this point. Can we evaluate that argument? Can't we? [inaudible discussion]. It's a bad one, isn't it? We can see that it's bad, we don't need to know anything.

'It's wrong to tell a lie.' 'Jane's telling her mum her hair looked good was a lie.' 'Therefore Jane's telling her mum her hair looked good was wrong.'

Can we evaluate that argument a priori?

Female: Yes.

Male: No.

Lecturer: I think we can. Is it a good one?

Group: Yes.

Lecturer: It's another good deductive argument:

'Tomato plants that have been fed well, kept warm and watered frequently usually thrive.' 'This tomato plant has been fed well, watered frequently but it's dead.' 'Therefore this tomato plant hasn't been fed properly.'

Female: That's inductive.

Lecturer: That's inductive.

Female: It might not have been kept warm.

Female: Because you want to know why it's dead.

Lecturer: Yes, why is it dead? Is it just that I pulled it out of its pot? Yes, we'd need to know a little more about that situation, wouldn't we? We'd need to have some background information.

What about this one? This looks familiar:

'If this liquid is acidic, it will turn litmus paper blue.' 'This liquid turns litmus paper blue.' 'Therefore this liquid is acidic.'

[inaudible discussion]. [I know, I know, acid turns litmus paper *red* not blue! MT]

It's good, it's a deductive argument.

'The last two springs were hot and sunny but the summers were awful.' 'This spring was hot and sunny.' 'Therefore this summer will be awful.'

Male: That's pretty true.

Female: It's not a priori.

Lecturer: Do we need to bring to bear background information?

Group: Yes.

Lecturer: What sort of background information would we need to bring to bear?

Female: Weather forecasts.

Male: We need more than two years' worth of information.

Lecturer: Yes, we'd want a lot more than two years, wouldn't we, in order to generate that.

Female: Marianne. Going back to 5. If the liquid is something other than acidic, it might turn the litmus paper blue too. Or are you bringing to bear some background knowledge?

Female: To be honest, when I was as school it was red ..

Lecturer: Let me have a look at that. I've got to the point where it's actually difficult to ...

That's a bad deductive argument. No, you're quite right. Did we say it was a good one?

Group: Yes.

Lecturer: Yes, that's because the last one that talked about litmus paper and acidic... No, that's a bad deductive argument. Can you see that? 'If P then Q', 'P' 'Q' therefore 'P.' Which is not 'If P then Q', 'P' therefore 'Q'. Do you see? You've got to be careful about that.

That's a fallacy of affirming the consequent. We'll talk about fallacies later on.

Male: I'm not sure that 3 is a good deductive argument.

Lecturer: 'It's wrong to tell a lie.' 'Jane's telling her mum her hair looked good was a lie.' 'Therefore Jane...

It is. [inaudible discussion]. Let me just see if it is.

'P' is, 'It's wrong to tell a lie.' Oh ... actually I'm going to have to go to the predicate calculus, excuse me a second.

Here's my interpretation: 'F of x', 'x is a lie'. 'W of x.' 'x is wrong', and 's' (which is a thing, an action in this case), 'Jane's telling her mum her hair looked good'.

So we've got 'all x', 'if x is lie then x is wrong'. That's the first premise.

L s, s is a lie. Do you see we've got that? Therefore, W s. Both of those are 's's. So 'everything is such that if it's a lie it's wrong'. Do you see that's the first premise?

Group: Yes.

Lecturer: 's is a lie', 'Jane's telling her mum her hair looked good is a lie.' Therefore s is wrong, 'Jane's telling her mum her hair looked good is wrong.' Good deductive argument.

Male: I can see the logic.

Lecturer: Oh, good.

Male: But in reality, the premises are probably ...

Lecturer: But we're not evaluating [inaudible discussion] the argument at the moment. I'm not saying anything about whether the premises are true or false. 'If widgets are havena' is nonsense.

One there and one there.

Male: Would it make any difference to 5 if you put 'if and only if' at the beginning?

Lecturer: It would always make a difference to change an 'if' to an 'if and only if.' Always. I found that students when they first start formalising something will immediately reach for a double arrow which is the 'if and only if' when actually all there is is an 'if'. Or all there is is an 'only if' and actually 'if and only if' is very different to 'if' and it's different also to 'only if'. So yes, it would make a huge difference.

Male: Would that then make it deductively valid, or a good argument?

Lecturer: It's exactly the same as one we had earlier when somebody was talking about 'if and only ...' If you have 'if and only if this liquid is acidic it will turn litmus paper blue,' yes, it would be.

Male: Yes, I thought it would. Okay. You...

Male: That sorted the question I was going to ask.

Lecturer:

Right, that's good, two for the price of one. (Slide 60) Oh, exercise to do at home already. Right, in that case, we can wait a bit for that because I'm sure there must be more questions.

I usually summarise. Have I not done that this week? Yes, here we go. So there are your exercises to do at home and I will give you the answers to these exercises on your answer sheet next week, so you can have a look at those at home.

(Slide 61) This week we've learnt that critical reasoning is normative, not descriptive. Do you remember that something's normative if it involves standards: either moral standards or rational standards – so right and wrong in some sense – it involves norms.

We've looked at the fact that there two types of following from. So there's deductive following from and inductive following from.

We've seen that deductive arguments are truth-preserving, or at least they are when they're good. Sorry, I didn't make that clear there but that should be clear there 'truth-preserving if good.'

They're also such that their being good is an either/or matter: do you remember they're monotonic? Once it's good, it can't be changed to anything else.

Deductive arguments are such that we can determine a priori whether they're good or not.

Inductive arguments are never truth-preserving; no matter how good they are, they don't preserve the truth in that very definite definition of truth-preserving that I gave at the beginning.

They're also such that their being good is a matter of degree. They're such that we can only determine whether they're good

by bringing to bear our experience, our understanding of the content and our background information.

(Slide 62) So that's it, folks, for today, but we've got time for questions. Sorry, somebody had a question.

Female: I've got a feeling that you mentioned earlier the 'all swans are white' business as being not deductive because somebody had found black swans. But then later you said that premises don't have to be true and that if you said, 'All swans are white, this is a swan, therefore it is white,' even if swans aren't white, surely that's still a good deductive argument, isn't it?

Lecturer: Are swans animals or birds? Let's pretend they're animals. They're birds, aren't they?

Male: They're both.

Lecturer: Here's an argument, and here's another argument. 'All swans are white.' 'The animal in the next room is a swan.' 'Therefore the animal in the next room is white.' Is that an inductive argument or a deductive argument?

Group: Deductive.

Lecturer: It's deductive. What about, 'Every swan I've ever seen has been white.' 'Therefore all swans are white'?

Group: Inductive.

Lecturer: Inductive. I could say, 'All swans I have seen have been white.'
It would still be inductive, wouldn't it?

Female: But the first one, the top one, the deductive one is correct even though it's not true in real life.

Lecturer: It is truth-preserving – that argument – even though one of its premises is false. We'll be doing this next week and it's something that people find very difficult. I think there's a sort of argument that goes inside people's heads that validity is good, truth is good, therefore validity and truth are the same thing. But they're not.

You can have a valid argument with a false conclusion and you can have an invalid argument with a true conclusion. So the notion of truth-preserving-ness is very different from truth *simpliciter* or *truth-generating*.

Male: Can I take you back to the argument about stabbing through the heart?

Lecturer: The plunging the dagger?

Male: Yes, because I would argue you're turning that into that argument with a hidden premise.

Lecturer: Do you remember where it was? I was talking about monotonicity, wasn't I? There we are.

Male: I would say there's a hidden premise there which turns it into a deductive argument.

Lecturer: No, every inductive argument you could find a hidden premise that turns it into a deductive argument. Taking that one: 'Every swan I've ever seen has been white.' 'I have seen every swan.' immediately becomes a deductive argument.

There's no problem with taking an inductive argument and adding a premise to turn into a deductive argument, we can always do that. I don't think I've done that here.

Male: I would suggest the premise is: Plunging daggers into the heart are enough to be fatal.

Lecturer: Not if you've already killed the chap from – so you die of an electric shock and then I plunge a dagger into your heart, I'm not guilty of murder. I may have had the intention to murder you, Christopher, but I didn't do it because you were already dead.

Female: It's ambiguous anyway because you don't know who does the 'his' refer to, the man, or Jones?

Lecturer: Yes, you would have to assume that there is somebody, we are assuming that there's somebody dead here, that what killed him was the dagger [inaudible discussion].

Male: But he's not dead because your policeman's trying to stop him.

Lecturer: He might have plunged it into his own heart of course as well. Yes. But we do sometimes just have to take it for granted that there is a person here that we're not introduced to.

I think in talking about the cat we talked about 'her', didn't we, without making it absolutely clear who we meant?

Female: But you've got three people there, you've got Smith the policeman, you've got Jones –

Lecturer: Smith and Jones and 'him', whoever he is.

Female: So if you've got three people, and it doesn't say who's heart it was, Smith or Jones.

Lecturer: No, it's true. Well, you've got the man who is 'him' clearly, I mean 'him' the man.

Female: It's the 'his' that is ambiguous because you've got three men there.

Lecturer: Jones killed the man, saw Jones plunge the dagger into his heart.

Male: It could be Jones's heart.

Female: Or Smith's.

Lecturer: I would think that any English person would read that case, it'd be an anaphoric reference back to this man here, I really do. I know you're all English but I think you're just saying it for fun.

One more.

Male: Supporting the point you made, about electrocution, you said he could have been electrocuted. But he can't because your policeman's dealing with someone who's still alive.

The policeman tried to stop Jones killing the man.

Male: And failed.

Lecturer: Sometimes you get daggers plunged into your hearts without killing you, can you? I don't know. Why are we talking about this? It doesn't seem to be in the slightest bit interesting.

Female: Are you saying overall that when we look as students, try to look at logic, deductive, inductive reasoning, that we basically take subjectivity out of the equation completely?

Lecturer: Not with inductive arguments because you've got to bring to bay your background knowledge on the argument. With *deduction*, your background knowledge can be just taken completely out of the equation. Yes.

Computers use deduction.

Female: They do, but with the induction, because it's subjective and however many we are here, we'd all come up with something, as we just have, something possibly different. Does that make it problematic?

Lecturer: Yes, there's a big difference between deduction and induction in that with deduction we have a very simple mechanical system that can be given to computers. Computers embody the mechanistic testing evaluation system for deductive arguments.

We can't do all deduction arguments that way. For example, if I say, 'Lying is wrong, therefore you shouldn't lie.' That looks like a deductive argument: it couldn't be the case that that's (the premise) true without that (the conclusion) being true. But we can't mechanistically test that because it seems to depend on the meaning of the word 'wrong'.

The propositional calculus and the predicate calculus are very powerful systems that we can use to test deduction, but they don't test *all* deduction, *deontic* logic – which is the logic of morality – it's still very much in its infancy. There's also *modal* logic which is less in its infancy by: if it is possible – that (symbol) just means 'possible' – then it is not...

I was trying to think of something more interesting but yes, you can see that's a deductive argument, but actually again, that seems to depend on the meaning of the word 'possible' there. So we might think that 'wrong' and 'possible' are both themselves logical words like 'and', 'or', 'if', but if so they act very differently and we don't know how they act.

But whereas we have a very good system for the vast majority of deductive arguments, if we can mechanise *probability* at all, it's only a very small part of it, so yes, there is a problem for mechanistic systems because you need to bring to bear all sorts of background knowledge. One final question.

Male: Is it correct, the speaker mentioned that it could be subjective – that inductive arguments could be subjective, is that correct?

Lecturer: I ban undergraduates from using the word 'subjective' until they're in their third year.

Female: I'm way past my third year!

Lecturer: You need to bring to bear some background information. To that extent, you're bringing your beliefs to bear on it.

A *subjective* state is a state of a subject accessible to conscious awareness, so perhaps subjective is not quite the right word. But it's certainly true that different people could be more or less inductively bold.

So some of you here might say, 'Every time I've seen Marianne, she's been wearing earrings, therefore next time I

see her she'll be wearing earrings,' even though you've only met me three times.

But some of you might wait to meet me ten times before you're prepared to say that or whatever.

Do you see? So some people are very inductively bold, they'll draw a generalised conclusion from very little evidence and other people want quite a lot of evidence before they'll draw an inductive conclusion.

So there's certainly an element of subjectivity in how inductively bold you are.

Right, we'd better stop there. The presentation for today is down here as are the arguments for the end of last week and today.

END AUDIO