

Lecture 4: Deductive Validity

Lecturer: Right, I'm told we can start. Hello everyone, and hello everyone on the podcast. This week we're going to do deductive validity. Last week we looked at all these things: have a quick look and remind yourself, because we're about to do a quick revision, as usual. So remind yourself what we did and then I'll ask you some questions. Enough of a reminder?

Who can tell me what it is for an argument to be 'truth-preserving'? Can anyone tell me what it is for an argument to be truth-preserving? I think I might have to give up these little revision sessions. (Laughter)

Female: The conclusion follows on from the premises?

Lecturer: The conclusion follows from the premises certainly, but that's not sufficient for an argument to be truth preserving, because that's true in inductive arguments as well deductive arguments. It's only good deductive arguments that are truth preserving.

Female: As it could only be that conclusion following on from those premises so there's no option.

Lecturer: There's no option.

Male: If the premises are true, the conclusion must be true.

Lecturer: Good, if the premises are true the conclusion *must* be true. So I see what you mean about no option but...

Male: [inaudible]

Lecturer: Yes, that was canonical; that was very good. An argument is truth preserving if there is no logically possible situation in which the premises are true and the conclusion false. So if the premise is true, the conclusion *must* be true – that logic says it must be true – there's no option as you put it. That's what it is for an argument to be truth preserving.

Do all deductive arguments preserve the truth?

Female: No.

Lecturer: Which ones don't?

Male: Bad ones.

Lecturer: The bad ones – exactly so. It's only good deductive arguments that preserve the truth, bad ones – let's see, do we ask this? No, okay, what do the premises of a *bad* deductive argument tell us about the conclusion? Can anyone tell me that? What do the premises of a bad deductive argument tell us about the conclusion?

Male: It's false.

Female: It's wrong. Well, it may not be true.

Lecturer: Hang on: 'false', 'wrong', go on – have another go.

Male: The conclusion doesn't necessarily follow from the premise.

Lecturer: The conclusion doesn't follow from the premises *at all*; the premises of a bad deductive argument are no reason whatsoever to believe the conclusion. That doesn't mean that the conclusion is false, the conclusion of a bad deductive argument can be true as we'll see later today. If by 'wrong' you meant 'false', the same thing applies.

Why aren't inductive arguments truth preserving? Can anyone tell me?

Female: Because they're can generate the truth.

Lecturer: I'm not sure quite what you mean by they 'can generate' another truth.

Female: They're not preserving the truth, they produce one.

Lecturer: No, that makes inductive arguments sound as if they can generate truth, and as a matter of fact, no that's not the case.

Male: It's because they're matters of degree.

Lecturer: That's not why. Inductive arguments are not truth preserving because it's always possible for the premise to be true and the conclusion false in an inductive argument, even the *strongest* inductive argument. Do you remember I said, "In the history of the universe the sun has risen every single day, therefore the sun will rise tomorrow."? That's a very strong inductive argument, but the sun might blow up tonight, we can't be certain that the sun is going to rise tomorrow. Logic has nothing to say about whether the sun is going to rise tomorrow just because it's risen every single day in the history of the universe.

What is it for an argument to be monotonic? A prize for anyone who can answer this, gentleman down there ... Sorry that's not fair: I'm picking on you but that's only because you've got every other answer right.

Male: No possible ...

Lecturer: Bad start! Can anyone remember?

Male: If the answer is 'yes' or 'no', there's no question of degree.

This is incorrect! It is only deductive *validity* that is monotonic, not invalidity.

You *can* make a bad deductive argument good by adding further premises.

Marianne Talbot

Yes, you're right. Deductive arguments are monotonic, which means that they're conclusive. You can add anything you like to a good deductive argument and you won't make it bad, and you can add anything you like to a bad deductive argument and you won't make it good. Whereas with an inductive argument, you can add premises and change a strong argument into a weak one and a weak argument into a strong one.

Male:

Last week, what we did is we had a very strong deductive argument on which you then added a premise which contradicted one of the premises.

I really struggled with – once you negated the other premise – that actually it didn't ...

Lecturer:

... Why that made a good argument?

Male:

Exactly, if everyone else got it then fine, but ...

Lecturer:

No I think I'm going to answer that question later on today. So if I don't answer it by the end of today perhaps you'll come back and ask me again.

Why is it always a matter of degree whether an inductive argument is good or bad? No? I'm embarrassing you aren't I?

That's because inductive arguments are *not* monotonic, you can always add something and you change whether they're good or bad. So they're either strong or weak, so whereas deduction is an either/or situation, induction is either strong or weak which is a matter of degree rather than either/or.

We'll go through the others quickly, what's the difference between a priori and a posteriori knowledge? Does anyone know?

Male: It's knowledge that you have now and knowledge that you may have in the future.

Lecturer: Not quite, but I see where you're going. A priori is knowledge you – go on ...

Female: Knowledge you had before having evidence for its existence.

Lecturer: Good, a priori is knowledge you have just from the ... so if I ask you whether there are any married bachelors, you don't have to leave your seat to tell me that there aren't. Similarly if I say, are there any two-sided triangles? You don't have to leave your seat to tell me that there aren't. This is a priori knowledge, it's knowledge that you have *without experience*, knowledge that you have just in virtue of grasping the concepts that I'm using. A posteriori knowledge is the opposite: you have to have experience; it's background knowledge about the world.

It's possible to evaluate deductive arguments a priori because they're, valid – good or bad – in virtue of their form and the logical words. Do you remember we reduced an argument to, "If P then Q. P therefore Q."? It's the 'if', 'then' and the 'therefore' that tells us whether the argument's good or bad, we don't need to know anything about background knowledge. Whereas with an inductive argument, as you'll see next week, we always have to know about the background knowledge of what we're talking

about: the *content* of the argument. Whereas deduction is topic neutral, I've just told you the answer to the last one. I could have asked you that and you could have had a quick win.

Okay, that's a quick bit of revision from last week. Any quick questions on that before we move on to today?

Female: Can you just tell me again please why inductive arguments aren't truth preserving?

Lecturer: Why inductive arguments aren't truth preserving? Well to be truth preserving is to be such that if the premise is true, the conclusion *must* be true. There's no logically possible situation in which the premise is true and the conclusion false. And an inductive argument is *never* like that because it relies on something called the 'Principle of the Uniformity of Nature'. We're always assuming in making an inductive argument that the future is going to be like the past. Logic says nothing about that; logic couldn't care less whether the future is like the past.

Female: So it's about content.

Lecturer: So it never guarantees truth, an inductive argument, even the strongest inductive argument, you never get a guarantee. We'll be looking at the evaluation of induction next week, so we'll learn a lot more about it then.

(Slide 4) Where we are now, this is the fourth lecture. We know what arguments are; we know how to analyse them and set them out logic book style; and we know the key characteristics of deductive and inductive arguments, and those exhaust the

types of arguments that there are. There's some disagreement about that, some people think that abductive arguments are another category, I'll discuss those a bit next week but I'm classifying arguments in such a way that there are only two categories.

This week we're going to learn how to evaluate deductive arguments. This is the bit that – in a way – you've all been waiting for. Because you've been dying to come in and tell me that premise is false or this conclusion doesn't follow or something like that. This is great, because what it shows is that you're all rational animals, you all know a good argument when you see one and you know a bad argument when you see one! What we're going to learn here is *what you know when you know that*.

Now we really can talk about true premises and conclusions and so on and the impact of that. Incidentally, just before I move on, those of you who noticed the fact that I had litmus paper turning acidic things blue last week rather than red: you're absolutely right, I had several emails and lots of people told me, you're quite right, acid turns litmus paper red. Actually I did know that, I don't know what happened but I didn't notice, so I do apologise, especially to the chemists amongst you.

(Slides 6 to 8) When a deductive argument is good, it's valid and its premises entail its conclusion, this is just a bit of terminology for you: a deductive argument, when it's good, it's valid and its premises entail the conclusion. So the relation of *entailment* is the relation between the premises and the conclusion of a valid deductive argument. When a deductive argument is bad it's invalid.

These are technical terms for logicians and indeed for philosophers generally. You probably use the word 'valid' as a sort of general 'feel good' term, and with that in mind you tend to

confuse validity and truth. Because truth is also a feel-good term isn't it? We want things to be true, we want things to be valid. So we tend to think that they're the same sort of thing. I hope today I'll get you out of thinking that – it's very important that you stop thinking that as quickly as possible. I'll remind you of it several times today.

A deductive argument when it's good is valid, when it preserves the truth it's valid, and its premises entail its conclusion. A deductive argument is valid and its premises entail its conclusion, when – and only when – there's no logically possible situation in which its premises are true and its conclusion false. In other words, when and only when it's truth preserving in the language we learnt last week.

A deductive argument is invalid whenever its premises fail to entail its conclusion. That's whenever there's a logically possible situation in which its premises are true and its conclusion false. I'm going over this terminology because when at home you're looking at your notes, you really need to learn these definitions, the definitions are *very important*. You can't learn logic properly without learning these definitions. So if you think vaguely, you're English and therefore you know what 'valid' means, you're going to go wrong almost straight away, I promise you. Do try and learn these definitions.

One way to determine whether deductive arguments are valid is simply to learn which argument *forms* are always valid and which argument forms are never valid. (Slide 10) So here you are; here's some argument forms that are always valid, modus ponens is 'if P then Q, P therefore Q. Modus tollens is 'if P then Q, not-Q therefore not-P' and so on, you can see all these. Here are some argument forms that are never valid, 'if P then Q, not-P therefore not-Q' do you want to have a quick look at those at

the moment and see if you can see why they're not valid, or they are valid?

Don't fear, I'm not going to make you learn them – I think that's the most boring way ever. Any questions about why those are valid and those are invalid? No? Okay, good.

So that's one way of evaluating deductive arguments is to know whether the form is a valid one or not. They could also be evaluated by means of truth tables, (Slide 11) I've got a truth table here, or tableau. I haven't got a tableau, but I could easily do one for you, in fact maybe I will just show you.

This looks like double-dutch and I remember when I was an undergraduate, I had the most *awful* trouble with truth tables, because the person teaching it didn't explain what was going on (sorry Mark if you're listening). What we've got here is a representation, if you like, of the truth conditions of both of the sentences that can constitute the argument and the impact of combining them in the way that we have.

So each of these is a *row*, I'm not going to ... well I *am* going to make you do one so you may as well listen, but I'm not going to ask you to rely on it. Each of these is a different row, and each row represents a different possible world or a different possible situation. So if we forget everything from this line, that way for the moment and just look at these. This is the world in which Q is true and P is true, and this is the world in which P is true and Q is false, and this is the world in which – tell me ...

Male/Female: P is false and Q is true.

Lecturer: You've got the idea. This is the world in which both P and A are false. Now you'll see that as we've got two, and we assume

bivalence when we're doing logic: bivalence is that every sentence has a truth value, so there are no sentences that don't have truth values. There are only two truth values, 'true' and 'false'. So a sentence can't be *neither* true nor false, nor can it be anything *other* than true or false.

Given that, there are only four possible combinations here aren't there? These are the four. One you've got that you only need to know what the import for truth is of the logical words 'if' and 'then' and I'm afraid you'll just have to take it from me that that's what it is. So that's the truth table definition of 'if P then Q'. P you'll see has exactly the truth values here that it has here, because if that's the world where P is true then P's got to be true there hasn't it? If that's the world where Q is false, then Q's got to be false there. So those two lines are exactly what comes from there.

Now you ask yourself, is this argument *truth preserving*? In other words, is there any possible situation in which the premises are true and the conclusion false? Well here's a situation where both premises are true but the conclusion is true, okay? So that's truth preserving. Here's a situation where it's not the case that both premises are true, so actually we don't even need to look at the outcomes. The same here, and the same here.

So do you see that in every logically possible situation where the premises are both true, the conclusion is also true? As this truth table exhausts all possible worlds we know that there is no logically possible situation in which the premises are true and the conclusion false. So that's a truth table test for a deductive argument.

I'm just going to make you do one for fun. I think I am anyway, yes I am, (Slide 12) I thought so. Let's ignore everything down from here to here just for the moment. Ah, now I shouldn't have

filled this in (laughter) it's just going to be on my head then isn't it? If P is true and Q is true, what's the truth value of 'P and Q'?

Male/Female: True.

Lecturer: Did you know that just from that? But you know that anyway don't you? If P is true and Q is true, it has to be true doesn't it? What about if P is true and Q is false?

Male/Female: False.

Lecturer: False, and what about that way?

Male/Female: False.

Lecturer: That one?

Male/Female: False.

Lecturer: Do you see where you get the truth value of 'P and Q' from these truth values here? So that's where those come from, and where does this come from?

Male/Female: [inaudible discussion].

Lecturer: It's just from the- that's right. Because that's true, true, false, false and this is true, true, false, false. Where does this come from?

Male: Q.

Lecturer: Q, true, false, true, false – same as over there. So is this a good argument or not? Tell me whether to put a tick or just a dash here.

Female: Dash. Oh no it's all false, yes.

Lecturer: A tick or a dash?

Male: Tick.

Lecturer: It's a dash, you're right. Why do you say that?

Female: Because it's all false.

Lecturer: No, that's not why you put a dash sorry, because that might be true as it is here. Do you see? What do you put here, a tick or a dash?

Female: Dash.

Lecturer: It's a dash; I've virtually told you that just now, haven't I? What about here?

Male/Female: Tick.

Lecturer: That's a tick yes, I won't tick the screen because I'll be in trouble. So what we've got here is a deductive argument: the premises on this side of this turnstile here (that's what it's called) and the conclusion on the other side. What we're looking for is a situation where all the premises are true and the conclusion false. If we can find that, we put a cross there and it's *not* a valid argument. But in this one, the only possible world in which both the premises are true is also a possible world in which the conclusion is true. Do you see how truth tables work?

As I say, we're not going to do those, I'll probably do them in another logic weekend next year or something. So keep an eye out for that if you're just dying to do truth tables. Tableau of course, you haven't met the joys of tableau yet.

Male: Could you just put some words onto an example for the first three columns?

Lecturer: The first three columns, yes I could make P 'it's Monday' and Q is 'Marianne's wearing jeans'. No, that won't do. I'm sorry I find it very difficult to think off the top of my head of an example. If I go back to the previous one I can give you the example that I just

had. 'If it's Monday then Marianne's wearing jeans', 'it is Monday', therefore 'Marianne's wearing jeans'.

Male: What's P and Q then?

Lecturer: So P is 'it's Monday' and Q is 'Marianne's wearing jeans'. So if it's Monday then Marianne's wearing jeans, it is Monday, therefore Marianne's wearing jeans. Okay? It would be very easy to think of one for that but I'm afraid when I'm in the middle of a lecture it's quite difficult to stop and think, oddly enough, but it is.

Both those methods involved learning a little bit about how to formalise arguments: we've used P and Q and so on. What we're doing is we're eliminating the English and replacing it with symbols. But what we're doing here is *informal* rather than *formal* logic, so you don't have to learn either of those systems. Instead we're going to learn the best way to evaluate a deductive argument informally, is firstly to set it up logic book style, and we know how to do that. (Slide 14) Secondly to construct the counterexample set, and thirdly to ask whether the sentences of the counter example set are consistent.

This is what we are going to learn. You know how to do that, you know how to identify the premises and the conclusion of an argument and set it out logic book style. All you need to know now is how to construct the counterexample set and how to ask whether the sentences of the counterexample set are consistent. That's what we're going to do.

The counterexample set (Slide 15) consists in the premises of the argument plus the negation of its conclusion. So the premises, just as they always were plus the negation of the

conclusion. A set of sentences is consistent – and this is another definition to take on board – if there's a logically possible situation in which all the sentences are true together. Notice a *logically* possible situation, not an *actual* situation or even an empirically possible situation, a logically possible situation.

Those are the definitions, now let's actually practice it (Slide 16) On the left hand side is an argument, a good deductive argument I think, yes. On the right hand side is the counterexample set of this argument. The counterexample set, you'll remember from the previous slide is the premises of the argument plus the negation of the conclusion.

Do you see, I have an argument here: "Deepak is a banker, all bankers are rich therefore Deepak is rich."? So, the 'therefore' marks the conclusion. Here, (on the RHS) I no longer have an argument, I have a counterexample set, a set of sentences that's all it is, it's not an argument because we're not saying that the conclusion follows from the premises here. We're just saying, "Here are three sentences, one of which is the negation of that conclusion". Do you see that? I've just tacked the words 'it is not the case that'.

Very importantly, you can't negate the conclusion by saying, "Deepak is not rich." That doesn't negate the possibility that Deepak doesn't exist, do you see what I mean? Whereas that ['it is not the case that'] does negate that possibility as well, so it negates everything, that's the negation of the conclusion.

So do you see how to make a counterexample set? That's not rocket science! Here's another example, (Slide 17) and another counterexample set: have a look at those and see if there's anything that ...

Female: If you met somebody saying that, you'd just say, "Well that's irrelevant", because you've brought something in that doesn't really relate ...

Lecturer: This?

Female: Yes – it doesn't relate to ...

Lecturer: Well at this point you're only half way through how this is used to test an argument. So remember what we're trying to do now is learn a pretty mechanical way of testing an argument. Let's say you found yourself, you've been reading the leader of your newspaper and you're really bothered about this issue. Who's going to win the next American election? I know you're all bothered about that. The newspaper gives you an argument for why it should be Mitt Romney or something like that, and you want to know what this argument is because actually you would have gone for Obama, so why are they arguing that?

It's a very complicated ... You identify the argument, having got the argument you still can't work it out, it's still too complicated to work out how to do it. This gives you a completely mechanical way to test whether an argument is a good one or a bad one. So nobody's using that at the moment to test an argument, all we're doing at the moment is creating a counterexample set, we haven't gone in for any evaluation at the moment.

Female: Is it that one should determine whether the end result was consistent?

Lecturer: We haven't got to that yet, so I said there are two things you've got to be able to do: the first one is create the counterexample set and the second one is to decide whether it's consistent. All we're doing at the moment is creating the counterexample set, I'll look at consistency in a moment.

Female: Can you change that from it is not the case then, could you say, "*However* Deepak is not rich."? "Is not a banker."? Sorry.

Lecturer: No.

Female: It's not the case then, have you got to-

Lecturer: No you're wanting to put 'however' in there because you're wanting another argument. That's an argument, (on the LHS) that isn't an argument. That's an argument, (LHS) that's the counterexample set (RHS) for the argument. So there's no 'however' or 'therefore' or anything in this one at all. Do you remember what I said about you've got to use, "It's not the case that Deepak is a banker," rather than, "Deepak is not a banker."? Because "Deepak is not a banker." is not a negation of, "Deepak is a banker." Or rich or something. This is a bad one isn't it?

Any questions about creating a counterexample set? You all look very tense and worried. (Slide 18) All of you provide a counterexample set for ... taking this right down the middle of the room there: this side start at the top and this side start at the bottom. So from the bottom up, and from the top down, just so

that we get some chance of getting them all done. Can you see if you can find a counterexample set?

[pause]

Anyone willing to have a go at number one? We'll wait a bit longer until somebody else ... Anyone willing to have a go at number four? Not yet, a bit longer?

Female: Number one-

Lecturer: Not yet, we'll just wait until there's a sort of critical mass of people.

Male: For number four will you just say ...?

Lecturer: I'm waiting until a few more people have been able to do it for themselves before we do it together.

Female: Is number three true? I'm sorry that's an entirely personal question, but is it true?

Lecturer: Number three is an argument, so it can't be true.

Female: I meant in the real world. (Laughter)

Lecturer: I'm sorry, number three is an argument, therefore in the real world it can't be true or indeed in any world. Arguments are not true. Do you mean the premise?

Female: No it doesn't matter. (Laughter).

Lecturer: Right who's willing to have a go at number one? There were loads of you a minute ago. Is it Keith?

Male: Yes, 'it's not the case that Bill was caught cheating'.

Lecturer: No give me the whole ... sentence one ...

Male: 'If anyone is caught cheating they'll be sent down'.

Lecturer: 'If anyone is caught cheating they'll be sent down'. Sentence two...

Male: 'Bill was sent down'.

Lecturer: 'Bill was sent down', and sentence three. Notice I'm doing sentences because of course if we're doing a counterexample it is not an argument, it's just a set of sentences. Number three...

Male: 'Bill must have been caught cheating'.

Lecturer: No you're doing the counterexample here.

Female: It is the case.

Lecturer: It is not the case.

Male: 'It is not the case that Bill was' - yes that's what I was going to say.

Lecturer: '*It is not the case* that Bill has been cheating', (or 'It is not the case that Bill must have been cheating'). Good, that's the counterexample for that one. Okay, what about number four, who's willing to try that?

Female: 'If Higgins was born in Bristol then he's not Cockney'.

Lecturer: 'If Higgins was born in Bristol then Higgins is not Cockney'.

Female: 'Higgins is either Cockney or an impersonator'. 'It is not the case that Higgins was born in Bristol'.

Lecturer: Sentence three: 'it is not the case'. Good, well done, 'that Higgins was born in Bristol.' So in each case there we've got the premises just as they were, so if you see those, those are

exactly as they are there, and the conclusion exactly as it is but with 'it is not the case that' tacked on the front. Do you see that's exactly what I've got for that one, and that's exactly what I've got for that one? Can you now do a counterexample for number two on that side of the room and number three on this side of the room?

[pause]

Put up your hand when you've got one.

[pause]

Anyone got number three? Hand up so I can see, good, okay. Right, number two, anyone got a counterexample for number two? Christopher?

Male: 'Damien Hirst does not follow rules and conventions', is the second one.

Lecturer: No I'm looking for the counterexample set at the moment.

Male: Yes, I'm saying the second sentence is that Damien Hirst follows neither rules nor conventions.

Lecturer: No, can anyone give me the counterexample set?

Male: Do you want to write this whole thing down?

Lecturer: Not particularly but I will do, just because I might – go on tell me what the whole thing is. In fact, read it out because it makes it easier for me.

Male: It's not possible to assess the artist Damien Hirst because it would be possible to assess his art only if he were following rules and conventions.

Lecturer: 'Only if Damien Hirst was following rules and conventions'.

Male: Then, it is a fact that he follows rules and conventions. It is the case.

Lecturer: It is *not* the case that he follows neither rules nor conventions. Just be like computers, all you're doing is tacking 'it is not the case' on the front. I see exactly what you're doing, you're thinking, "There's a double negation here and I'm going to get rid of it." But I don't want you to do that. So 'it is not the case that he follows neither rules nor conventions'. Okay, good and number three.

Male: Sorry, could the Hurst question be split into two sentences?

Lecturer: Well I haven't.

Male: No, but in analysing it.

Male: It is not possible to assess the art of Damien Hirst, period. It *would* be possible to assess ...

Lecturer: Yes you could do that, because that definitely is an 'and' isn't it? It's not possible to do this *and* this is why, yes.

Male: Could you argue that the first sentence is the conclusion, that 'it is not possible to assess the art of Damien Hirst', is that the conclusion?

Male: That was what I was trying to argue.

Lecturer: I think that, it seems to me that ... hold on.

Male: You could say 'therefore it is not possible to assess the work of Damien Hirst'.

Lecturer: If I'm trying to think I can't do it... I think you're right actually, it looks as if that's the conclusion doesn't it? And that and that are the premises? Yes I wonder why I put that in like that. Shall we do that again actually because I think it's got rather confusing? So I'm going to make this two again, and let's do it as you rightly say there. Because I think you're right, that is a much better analysis of that argument. Let's set it out logic-book style first, and then provide the counterexample and then we can see what we're all doing. So logic-book style, we want premise one,

premise two and the conclusion. Premise one is.

Male: 'It's possible to assess the art of Damien Hirst, only if he were following rules and conventions'.

Lecturer: '...To assess the work of Damien Hirst only if he were following rules and conventions'. Premise two...

Male: 'He follows neither rules nor...'-

Lecturer: 'Damien Hirst follows neither rules nor conventions', and conclusion?

Male: Therefore it's ...

Lecturer: 'Therefore it is not possible to assess the art of Damien Hirst.' Notice we had to get rid of cross references and things? I'm doing that by putting Damien Hirst in rather than 'he'. Okay can we now have a counterexample of that?

Male: It is not the case that it's not possible to-

Lecturer: 'It is not the case that it is not possible to assess the work of Damien Hirst'.

- Female: It's full of double negatives: my language teacher would be turning in her grave!
- Lecturer: Yes but the thing is if you try and get rid of the double negatives as you're doing it you might introduce an error. That's what I don't want you to do, I just want you to create the counterexample. You're right, I should probably have not used so many double negatives. Sentence one there, its premise is the same as premise one. Sentence two there is the same as premise two.
- Now the one that you're interested in, number three ... That's pretty straight forward isn't it in terms of the argument? Let's have just the counterexample then.
- Class: It is not the case that [inaudible discussion].
- Lecturer: The counterexample is the whole thing, so 'if you live alone or with someone who is mentally ill you are treated as a single person for council tax purposes'. 'Jennifer pays council tax as a single person, therefore' – no we don't want 'therefore' because it's a counterexample – 'It is not the case that Jennifer lives alone or with a person who is mentally ill'. To answer your question which I think was, is that premise true?
- Female: It obviously wasn't, I was being a bit cheeky, I wanted to know if it's fact.
- Lecturer: This first premise is true.

Female: It actually is?

Lecturer: When my mum, who had Alzheimer's, was living with me I was treated as a single person for council tax purposes. You should be too if you're in the same situation. Just let me go back though, you asked, "Is that true?" That is an argument and it can't be true. Do you remember back in the first week I said *sentences* can be true or false but arguments can only be good or bad, or valid or invalid, or strong or weak? So when you said, "Is that true?", you couldn't have meant the – but I knew you meant one of the premises, and it probably had to be that one. But yes, indeed, that is true.

Have you got the idea of counterexamples? All you're doing is negating the conclusion of the argument and setting out a set of sentences. The first of which are premise, premise, premise, and a negation of the conclusion. That's the counterexample set. That's funny, I wasn't expecting that.

[pause]

Nevertheless that's what it is. Okay, what I was going to go on to say and have clearly forgotten to put on the slides is, are these counterexample sets consistent? Because that's the next step: until you've decided whether they're consistent or not you haven't been able to determine whether the argument's a good one. Given that a set of sentences is consistent only when there is a possible situation when they're all true together, let's have a look at each of these and see whether we can find such a situation.

Now you can't read my writing can you? Interestingly enough neither can I. 'If anyone is caught cheating they'll be sent down',

'Bill was sent down', 'it is not the case that Bill must have been cheating'. Is that set of sentences consistent?

Female: Yes.

Lecturer: Could they all be true together?

Male/Female: Yes.

Lecturer: Could somebody give me a situation in which they could all be true together?

Female: Bill may have been sent down for another reason other than cheating.

Lecturer: Yes, exactly, he could have been sent down for another reason. So those premises could be true and the negation of the conclusion could be true as well. What does that tell us?

Male: That there isn't a ... it's not a valid argument.

Class: [inaudible discussion]

Lecturer: There are some very interesting answers coming in here. Let me ask the question again because I hoped in asking the question

to make it clear what I was getting at. The premises of this argument can be true together and the negation of the conclusion of the argument can be true at the same time as both of the premises are true. What does that tell us?

Male: It's not a valid argument.

Lecturer: It tells us that the set of sentences of the counterexample set is consistent, therefore it's a bad argument. It's an invalid argument. Remember ... I need someone to come up and hold a piece of paper for me, don't all jump at once ... Would you like to come up and hold the piece of paper for me?

Female: If I'm tall enough.

Lecturer: I don't think you need to be tall.

[Pause]

Now don't tell me you can't read that, I used to tell my undergraduates that if they can't read my writing by the end of term they were going to fail their exam. Would you like to stand just a bit further over there, because I think the light ...? Can you read that? I'll read it out to you, an argument is valid- actually you might have to sit down and hold it up there so it's out of the light.

An argument is valid if, and only if there is no logically possible situation in which the premises are true and the conclusion false. So an argument is truth-preserving – i.e. valid – if it matches those conditions, if it satisfies those conditions. If

you've got a counterexample set, what you've got is a set of sentences that is the premises and the negation of the conclusion haven't you? So it's as if we're saying, but what if this conclusion were false? Is there a possible situation in which the premises are true and this negation of the conclusion is also true? Because if there is then it can't be valid can it? Do you see where I'm coming from?

Where were we?

Male: So are you testing the argument?

Lecturer: Yes, that's exactly what we're doing. So there's the counterexample set for argument one, do you remember? 'If anyone is caught cheating they'll be sent down', that's premise one on there. 'Bill was sent down', that's premise two. We've negated the conclusion, okay? So here we have the counterexample set. If this counterexample set is *consistent* then it represents a possible situation in which the premises are true and so is the negation of the conclusion.

Female: [inaudible]

Lecturer: In other words a situation in which the premises are true and the conclusion is false. If the counterexample set is consistent, then the argument is *invalid*. Do you see why you need to make the difference between truth and validity and consistency and so on? They all interact with possible worlds so you've got the definition of 'validity' there and in order to mechanically test

whether an argument is a good one, all we have to do is negate the conclusion and just see if the set is consistent.

If we can think of a situation in which all those three are true together, then we have a counterexample to the argument. Do you see?

Male: Yes, clever.

Lecturer: So shall we do another one? Thank you very much, you can have that. (Laughter) Let's do another one, 'if Higgins was born in Bristol then Higgins is not Cockney'. 'Higgins is either Cockney or an impersonator'. 'It's not the case that Higgins was born in Bristol'. Is that set consistent? Is there a possible situation where all those sentences are true together?

Male: Yes.

Lecturer: Can anyone tell me what it is?

Male: Higgins is an imposter?

Lecturer: Or an *impersonator* isn't it? Not imposter, impersonator. Okay, he's an impersonator, that would make all those true wouldn't it? Immediately we not only know that our argument is bad, we also know *why* it's bad, we've got a counterexample to the argument. We can say, "No that's not a good argument because if Higgins is an impersonator, then Higgins *wasn't* born in Bristol. Your

premises are true and yet your conclusion's false." Are you with me? Do you see how it works?

We'll just do the harder ones. That's the one we messed up. 'It would be possible to assess the work of Damien Hirst only if he were following rules and conventions'. 'Damien Hirst follows neither rules nor conventions', and 'It's not the case that it's not possible ...' we do want to get rid of that double negation don't we? 'It is possible to assess the work of Damien Hirst'. Could those all be true together?

Male: Yes.

Male: No.

Male: There's an 'if and only if' isn't there?

Lecturer: Is there?

Male: 'Only if', does that mean the same as -

Lecturer: There's an 'only if' but not an 'if and only if'. An 'if and only if' is different from an 'only if' because it's got an 'if and' in front of it.

Female: You can't get rid of those negations can you? 'It is not the case that it is possible to assess the art of Damien Hirst'.

Lecturer: 'It is not the case that it's not possible to assess the work'... I think if you say, "It's not the case that it's not possible to assess the work of Damien Hirst," that is the same as, "It is possible to assess the work of Damien Hirst," isn't it?

Female: Yes.

Lecturer: Let me write these out again, because I messed that up. So sentence one is, let me just write this out: 'It would be possible to assess the work of Damien Hirst only if he were following rules and conventions'. Sentence two is, 'Damien Hirst follows neither rules nor conventions.' The negation of the conclusion, having taken out the double negation is, 'it is possible to assess the work of Damien Hirst'.

So the question is, that's the counterexample set and we're asking is it consistent?

Male: No.

Lecturer: What do we know then? The counterexample set is *inconsistent*, in other words there is no possible situation in which the premises and the negation of the conclusion are true together. So the argument is *valid*, that's right.

Do you see you've got to be careful of another hiccup here? Because if the counterexample set is *inconsistent*, then the argument is *valid*, and if the counterexample set is consistent then the argument is invalid. Again, don't go – consistency is good, validity is good, therefore a consistent counterexample set

means a valid argument because do you see that'll get it wrong; you'll get your possibilities the wrong way. Is that a question?

Male: I'm struggling with that, I can see that being the case if it was an 'if and only if'.

Lecturer: But it's not.

Male: So why can't there be another case where you can assess the work of Damien Hirst?

Lecturer: Let's just have a look. Yes it's a necessary condition of assessing the work of Damien Hirst that he's following rules and conventions. But he doesn't follow rules and conventions.

Male: It doesn't say it's the *only* condition.

Lecturer: No but it does say it's a *necessary* condition.

Female: In this particular example ...

Lecturer: No it isn't an 'if and only if'.

Female: In this particular example, what is the difference between saying 'it would be possible to assess his art only if he were following

rules and conventions', or I say 'it would be possible to assess his art if and only if he was following'-

Lecturer: Okay, I think this is where I've been at it too long, because I know that people have huge problems with 'if and only if'. (Flip chart) So there's 'if, then', there's 'only if' and there's 'if and only if'. Those are the three things that people tend to get confused, can you all see? 'If P then Q', 'P only if Q', and 'P if, and only if Q'. If we say 'if P then Q' we're saying that P is sufficient for Q aren't we? If we're saying 'P only if Q', we're saying that Q is necessary for P. If we say 'P if, and only if Q' we're saying P is necessary *and* sufficient for Q. So we're adding that ('if P then Q') to that ('P only if Q') to get that ('P if and only if Q').

Female: If I just go to Damien Hirst, if you go back to the previous to save you writing out again, 'it would be possible to assess the work of Damien Hirst only if he were following rules and conventions', it's true.

Lecturer: So that says 'it's a *necessary* condition of assessing the work of Damien Hirst that he follows rules and conventions'. Let me write that down because I think that might ... 'it is a necessary condition of assessing the work of Damien Hirst that he follows rules and conventions', is that right? 'Damien Hirst doesn't follow rules and conventions': sorry, this is the counterexample set isn't it?

Female: It's not possible.

Lecturer: Well no it is possible to assess isn't it? 'It is possible to assess the work of Damien Hirst.

Female: Suppose you go straight to the bottom, you can assess Damien Hirst's work on another premise other than rules and conventions.'

Lecturer: No, but this says it doesn't matter what else might be – what other way ... if it's a *necessary* condition that he follows rules and conventions, that means if he *doesn't* follow rules and conventions it isn't possible to assess his work. That's inconsistent with that. That set is inconsistent.

Female: That's why I asked the question about [inaudible].

Lecturer: But do you see -

Female: No I don't, but perhaps we can take it implied because if the only way to assess – that's why I thought it would be an 'if and only if'. If and only if ...

Lecturer: If it was an 'if and only if' it would be ...

Female: Rules and conventions to assess his work.

Lecturer:

If it were an 'if and only if' then the argument might be okay. But it isn't an 'if and only if', and it's *this* argument we're evaluating, not a different argument that had an 'if and only if' in it. Do you see what I mean? You absolutely mustn't ... if you're analysing your argument set out logic book style. If that's the argument you want to analyse, you must analyse *that* argument not another argument that you think might work. Because then you're creating another argument.

Do you see here? This isn't an 'if and only if', it is an 'only if', if it were an 'if and only if' then you would have, "It's a necessary and sufficient condition of assessing the work of..." Actually it still wouldn't, it doesn't change it, in this case. 'If and only if' and 'only if' are very different in other cases though; you've got to be very careful.

So that one's valid, that one's valid – no, that one's invalid and that one's valid, is that right? We haven't yet looked at the other one. Let's have a quick look at that ... Do you see how if you can get used to doing all this, you will think much more clearly than you ever have before? You will make the difference immediately between 'ifs' and 'only ifs' and 'if and only ifs' and so on, and you'll see how possible worlds interact with truth and with possible truth in such a way that will be really helpful for you.

Here we are, 'If you live alone or with someone who's mentally ill, you are treated as a single person for the purposes of council tax'. 'Jennifer pays council tax as a single person'; 'It's not the case that Jennifer lives alone or with a person who's mentally ill'. Is that set consistent?

Male:

Yes.

Female: No.

Lecturer: Half of you are saying yes and half of you are saying no, put up your hand if you think, yes it's consistent. Sorry, put up your hand if you think, no, it's inconsistent. Okay, give me the counterexample then, sorry have I got that wrong? You think it's inconsistent, therefore ... okay, those of you who think it's *consistent*, give me the counterexample.

Female: That is the counter example.

Lecturer: Sorry, yes that is the counterexample, what I mean is the counterexample to the argument. If those are consistent then we have a counterexample to the argument don't we? Are you with me? So what is that counterexample? Those who think that this set is consistent, they believe that there's a possible situation in which all of these are true together. If you tell me what that possible situation is, and it really is a possible situation and we agree, then we will have found a counterexample to that argument.

Female: There are other reasons for being treated as a single person for the purposes of council tax other than - it doesn't say 'if and only if'.

Lecturer: Good, you're absolutely right. The fact is, if you're treated as a single person for the purposes of council tax if you have red hair, as well as being single or living with a person who's mentally ill – then if Jennifer has red hair, she could be being

treated as a single person couldn't she? So we see that this counterexample set is consistent, i.e. there is a possible situation in which the sentences, the premises and the negation of the conclusion are all true together. So the argument is *invalid*, you've got it, very good.

It's very tricky, but actually once you understand how the concepts work in together, it's not tricky at all, it'll become very easy for you. Everything's come off my chair, therefore I've got nowhere to stand!

Female: Just before we ... just on the question of terms: a set of sentences like that is a counterexample?

Lecturer: A set of sentences like this [the premises plus the negation of the conclusion] is a counterexample *set*.

Female: Then the counterexample is ...?

Lecturer: The situation in which those sentences are all true together. Let me make that clear. First you have the argument; then you create the counterexample set by negating the conclusion of the argument; then you ask yourself if there's a situation in which the sentences are consistent. If there is, you have a counterexample to the argument. A situation in which the premises are true and the conclusion false and therefore the argument is *invalid*.

Shall I say that again?

Male: Yes please.

Lecturer: Here we have an argument: somebody is asserting this sentence [the conclusion] on the grounds of these sentences [these premises], that's what an argument is. We want to know whether the argument is a good one or not, so we set it out logic-book style, we identify the first premise and the second premise and the conclusion. Then we want to evaluate the argument, so we take the conclusion of the argument and we tack 'it is not the case that' on the front of it, thereby making the counterexample set for the argument.

We then stand back and say, "Okay, are the sentences of the counterexample set consistent, *could* they all be true together? Is there a situation I can imagine in which they're all true together?" If there is, then you say what that situation is and *that* is the situation in which that argument is invalid. Because what you've shown is that this argument is not truth-preserving: the premises can be true and the conclusion false, because the premises can be true *together with the negation of the conclusion*, therefore the argument is not truth-preserving.

Male: That seems slightly long-winded to me, if you take example one, I would say 'if P then Q' then it says 'Q therefore P' which on one of your earlier slides you've already shown is not a valid argument.

Lecturer: You're absolutely right, another way of evaluating arguments is to determine *the form* [the structure of the argument]. But what I want you to be able to do is to go away from here and informally evaluate arguments that you come across. It would be very nice

if all the arguments you come across could have the form of modus ponens which is the one that you gave, but I'm afraid they won't have. What's more they won't be easy to put into P and Q form necessarily either, and you don't know how to do that.

So what I'm doing is giving you a way that - it's not going to be easy to start off with to learn how to apply this, but having learnt how to apply it, you can apply it to any deductive argument at all. It won't always be easy because as you can see sometimes it's quite difficult to see whether a set of sentences is consistent or not. But you've got a way of going into even the most complicated argument by doing this.

That's a way of evaluating whether or not a deductive argument is valid. For a logician, it's a sufficient condition of a deductive argument being good that it's valid. So if we're logicians now, we know exactly which of the arguments on the previous slide is a good argument and which isn't. But in everyday life we want more of a deductive argument than it simply being valid before we count it as good, let me illustrate that for you.

Are these arguments valid? (Slide 20)

[pause]

I'm going to send you away with your heads spinning today.

[pause]

Put up your hands if you think, 'yes'... Okay, let me tell you: they're both valid. Let me explain why. An argument is valid if there's no possible situation in which the premises are true *and* the conclusion false. In this one (LHS) there's no possible situation in which the conclusion is false is there? So how could there be a possible situation in which the premise is true *and* the conclusion false? There couldn't be could there?

Male: Grass isn't always green!

Lecturer: That's irrelevant because what matters here is that as there's no possible situation in which the conclusion is false, you know immediately it doesn't matter what the premise is, that there is no possible situation in which the premises are true *and* the conclusion false. Do you see? This one (RHS) is the mirror image of that: if there's no possible situation in which the premises are *true*, then how could there be a possible situation in which the premises are true and the conclusion false? There couldn't be. And therefore both those arguments are valid.

But they're not *good* are they? I mean, no one would be happy with an argument like that, and that's why an argument must have *more* than validity if it's going to be a good argument.

Female: Because I was curious, do you call something an argument, if it's just a few sentences stuck together – grass is green, and then you say, "therefore ..."?

Lecturer: Well, anything can be an argument. If you're making that claim [that the conclusion of the LHS follows from the premises]: actually a logician would make this claim quite often because we know that from contradiction anything follows, you might say, and if something's a necessary truth, it follows from everything. So you could argue like that, but you wouldn't use it with somebody who didn't know about- these are the 'paradoxes of entailment', they're called.

Those arguments are bad because of the paradoxes of entailment. If a deductive argument has contradictory premises then it's *always* valid and that's because there's no possibility of its premises being true, and therefore no possibility of its premises being true *and* its conclusion false. Do you see why it's valid?

Male: Yes.

Lecturer: That it satisfies the definition. This one, if an argument has a tautological conclusion then it's always valid – again – because there's no logical possibility of the conclusions being false, so no possibility either of its premises being true *and* its conclusion being false. These are paradoxes because they're very irritating. We would rather the definition of validity didn't generate paradoxes like this, but it's a bit like - whales for example are fish aren't they? They're obviously fish, except they're not, they're mammals. Why are they mammals? Because they have live young.

It's an unfortunate thing that the way we define things – whales turn out to be mammals rather than fish – well we could just change all our definitions and so on. But actually for all other purposes the definitions work perfectly well, so it's exactly the same here.

With this definition of validity that I've given you, we can make computers do all sorts of things, get to the moon, do all sorts of things. Let someone jump from a spaceship goodness knows how many thousand feet or whatever it is. How far was it?

Female: Twenty-two thousand!

Lecturer:

Anyway, we like the definition but can you see that again you've got to understand these definitions; because once you understand the definitions you will see why the paradoxes of entailment are valid. They show us that validity isn't enough to make an argument good for everyday purposes; we also need the premises to be *relevant* to the conclusion.

So coming back to your question here, can you make anything an argument just by putting 'therefore' in front? The answer is, it depends how you use it. For example, almost anything can be relevant to almost anything else given the right context, and I think ...here's an answer: 'the sea is salt is true, therefore Melbourne is in Australia.' You might think that couldn't possibly be an argument, they're not relevant to each other.

(Slide 24) Then we say that we're in the middle of a game show and we've been told that if the first sentence is true then the second sentence is also true. The first sentence happens to be, 'the sea is salt' is true, the second sentence is 'Melbourne is in Australia'. We're very thick, we didn't know either of these things, so we look up 'the sea is salt' and we find out 'the sea is salt' is true, *therefore* Melbourne's in Australia. Do you see how it does become an argument? All I've done is created a context.

The thing about human beings is we can use arguments; we can use the language in all sorts of ways. Almost anything can be relevant to almost anything else. There's a type of logic called relevance logic that is trying to work out the logic of when something is and isn't relevant – fascinating stuff but you really don't want to think about that right now.

(Slide 25) Context is very important in the evaluation of argument, if my sat-nav tells me to turn left one minute and then

turn right the next ... actually it probably is being inconsistent knowing my sat-nav. But it's not being inconsistent is it?

Male: No.

Lecturer: Because the context is different, and I've just put the definition of inconsistent in there. They're inconsistent if the set can't all be true together, and actually it's important that they can't all be false together, I should have put that in but never mind.

We've built in context into the way we've talked about arguments, so if we mentioned Sue in one paragraph and Sue in the next paragraph, we're assuming that the context is such that the Sue in each case is the same. So do you see, we've built in context? If I say, "I'm hungry." And you say, "I'm not hungry." it's contradictory: it's neither true nor false.

When you're evaluating an argument you need to be aware of the context in which the argument is being used. You almost certainly will be because in informal logic we very rarely take an argument out of its context so that shouldn't be a problem for you. But in these lectures we've been assuming that contexts are always constant.

Another thing we need in addition to validity is 'soundness'. (Slide 28) The fact that an argument is valid doesn't mean either that its premise must be true or that its conclusion must be true. That's bothered some of you as we've gone through these lectures, you've said, "But that's not true," sometimes, and I've always said to you, that at the moment we're not evaluating arguments.

When we're evaluating arguments the question of the truth of the premises becomes important. A valid argument can have

false premises and a false conclusion and here's an example... (Slide 30) Do you see that that's valid? It's truth-preserving isn't it? If *that* were true and *that* were true, then *that* would have to be true wouldn't it? So this is a valid argument but the premises are false and the conclusion is false. It's valid because if its premises *were* true it would be logically impossible for its conclusion to be false, okay? So truth-preserving doesn't mean that either the premises or the conclusion are *actually* true.

Here's another one, (Slide 31) a valid argument can have some true premises, some false premises and a false conclusion. Have another look; do you accept that that's valid?

Male/Female: Yes.

Lecturer: You do? Some of you say no? They're not prepared to admit it even if they did. That's again valid because if the premises *were* both true it would again be logically possible for the conclusion to be false. So again the argument preserves truth, even though you've got one – I think that's a true premise, but that's certainly a false premise and that's a false conclusion.

(Slide 32) And finally a valid argument can have false premises and a true conclusion: have a look; do you accept that that's valid?

Male/Female: Yes.

Lecturer: No? It's valid because if that [premise one] were true and that [premise two] were true, that would have to be true wouldn't it? Okay so it's valid, but that's false [premise one], that's false

[premise two] and that's true [conclusion], yes? So again it's valid because if its premises were true then its conclusion would have to be true. This is where you distinguish truth preserving-ness or validity from truth itself.

If a deductive argument is valid this is going to be *always* because its premises don't entail its conclusion i.e. because it doesn't preserve the truth. In that it's logically possible for its premises to be true and its conclusion false. That doesn't mean an invalid argument must have false premises or a false conclusion, and here it's an invalid argument with false premises and a true conclusion. (Slide 35) I've got past the point of being able to tell this. Is that an invalid argument?

Male: Yes.

Lecturer: Those could both be true [the premises] and yet that [the conclusion] false, right?

Male: Yes.

Lecturer: Yet that's true, that's false and that's true. Is that right? Can you see why it's invalid? Despite the way its truth conditions ... Here's an invalid argument with true premises and a true conclusion ... (Slide 36) You accept that?

Female: Yes.

Lecturer: You accept that both the premises and the conclusion are true? But that the argument does not preserve the truth? Here is a little table (Slide 37) that might help you: if an argument has true premises and a true conclusion, it could be either valid or invalid; if it has false premises and a true conclusion, it could be either valid or invalid; if it's got true premises and a false conclusion, it *must* be invalid, that's why it's useful. Because validity preserves truth, you cannot have an argument that's valid, and has *true* premises and a *false* conclusion.

Let's see if you can do this, we've got 10 minutes left. If an argument is invalid it will have a false conclusion. (Slide 38) Put your hand up if you think that's true [number 1]... Put your hand up if you think it's not true... [Applause]

Well done, it's not true. If an argument has true premises and a true conclusion it will be valid [number 2]?

Male/Female: True.

Male: Not necessarily.

Lecturer: You see, you're confusing validity and truth again, you're thinking, "If it has true premises and a true conclusion then it must be a good argument." Must it?

Male: No.

Lecturer: No, so you can have an argument with true premises and a true conclusion and that is invalid. Just to prove that we have one there I think, true premises, true conclusion and invalid, yes?

If the premises of an argument contradict each other, the argument will be invalid, true or false? [Number 3]

Male: False.

Female: True.

Lecturer: False, it must be valid mustn't it? This is one of the paradoxes of entailment. If the premises contradict each other, there is no possible situation in which the premises are true and therefore there isn't any possible situation in which the premises are true *and* the conclusion false. So the argument must be valid: that's a paradox. If you're having trouble understanding that, don't worry, it is difficult to make sense of that. But if you go over your notes later I hope that should become clear to you.

Finally an argument if it's valid is valid because its premises are true and its conclusion false [number 4].

Male: No.

Lecturer: Put your hand up if you think 'no'... Good, well done. I'm sure you'll all go home and pore over these until we come together again. These questions are very useful in testing your understanding of validity and truth and so on.

(Slide 39) If we're interested in truth of our conclusions as well as the validity of our arguments, and of course we're all interested in that. It's only logicians who are interested only in the validity of arguments. We want our arguments to be *sound* as well as valid, an argument is sound if it's valid and such that its premises are true.

Given the nature of validity, the conclusion of such an argument is logically guaranteed to be true. Here's a little table again, (Slide 41) if it has false premises and the argument is valid, it's unsound. If it's got true premises and it's an invalid argument it's unsound, it's unsound if it's invalid and false premises. It's only sound if you've got true premises and a valid argument because then the conclusion *must* be valid. So sound argument is what we're all looking for, all the time: that's what we really want.

(Slide 42) There are two questions you've got to ask of an argument if you want to know whether its conclusion's true, if you want to evaluate that argument. (i) Is the argument valid and (ii) are the premises true? To test this you use counter example sets and whether it's consistent or not. To test that, I'm afraid you've got to go out with your microscope and actually look at the world, we don't dirty our hands with things like that.

Though to a logician an argument's being valid suffices for its being good, in everyday life we also want the premises to be *relevant* to the conclusion, and we want the premises to be *true*. (Slide 44) But, what about this, is this a good argument? Is it valid?

Male:

Yes.

Lecturer:

Well done, you're absolutely right it is valid. The reason it's valid is how can it possibly be the case if that's true and that's false when they're exactly the same? Do you see? This is a circular argument and all circular arguments are valid; what's more they're valid, sound, and these premises couldn't be more relevant to the conclusion. But it's still not a good argument is it in terms of how we would think of an argument? Because it's not in the slightest bit persuasive.

So an argument is persuasive *only* if someone might accept the premises and yet deny the conclusion. That's why we want to use an argument because you say, "Not P." and I say, "Hang on if you accept this and this, you've got to accept that P is true." So I'm assuming that you're accepting the premises and yet want to deny the conclusion. If I can show the conclusion follows from the premises then I've made my case, you have to accept my argument.

But nobody is going to accept the premise of, 'all whales are mammals' and yet deny the conclusion 'all whales are mammals are they'? The argument is circular, all circular arguments are valid and relevant and many of them are sound. But thanks to the monotonicity of validity, such arguments can be persuasive because all I need to do is just add in lots of other premises, so that you don't notice that the conclusion is there amongst the premises.

Because you are all - you may not think it right now, but you are all *validity detectors*, that's what as rational animals you are. You hear that the argument is valid and you think, "Okay, that's a good argument." But actually it's only valid because it was circular, and you've got the conclusion hidden in there amongst the premises. The minute you've got the conclusion amongst the premises it doesn't matter what you add to the argument, you

won't stop it being valid because deduction is monotonic, (Slide 48) as we looked at last week.

So even if we're persuaded by circular arguments, we shouldn't be... (Slide 50). In order to evaluate a deductive argument we've got to ask the following questions: is it valid? That you test with counter examples and so on; is it sound? That's the one where you've got to take your microscope out and find out whether it's true, the premises; are the premises relevant to its conclusion? Finally, is it circular?

(Slide 51) Here are your exercises to do at home, and that's a list of what you've learnt this week if you're still able to keep your... I'm going to test you on it next week so make sure - these are the questions I'm going to ask you next week. That's it for today, okay.

[Applause]

END AUDIO