Lecture 5: Inductive Strength

Lecturer: Right, I think I'm going to start, because we're only going to do revision the first time. Though, I've taken pity on you this week, after the appalling performance last week, and I'm not asking quite so many questions.

(Slide 2) Last week, if you remember, we learnt how to evaluate deductive arguments as either valid or invalid, by constructing counterexamples. Can anyone tell me what a 'counterexample is?

Female: It's an argument which we've produced as exactly the premises of the argument, and a conclusion which means the opposite by a, 'not', 'It is *not* the case that…'

Lecturer: Now, there was only one thing wrong with what you said. So, what was your name?

Female: Orsa.

Lecturer: Orsa said a counterexample is an argument that has exactly the premises plus the conclusion with, 'not' in front of it. Now, there's something wrong with that. That was a good answer, apart from this…

Male: I think it's not an argument?

Lecturer: It's not an argument; it's very importantly not an argument. A counterexample is a set of sentences, so you make a counterexample from an argument, because if you remember, an argument is a set of sentences where one thing is being asserted on the basis of the others. Whereas, a counterexample set, none of them is being asserted, you're just checking to see …

So, what are you checking the counterexample set for? Can anyone tell me?

Male: Consistency.

Lecturer: You're checking it for consistency, and what is it for a set of sentences to be consistent? Can anyone tell me what it is for a set of sentences to be consistent?

Male: They all reflect the same thing? They are all the same.

Lecturer: No, not necessarily. They're almost certainly relevant to each other, so in that sense they deal with the same thing, but that's not what makes them consistent.

Female: Are they valid?

Lecturer: No, sentences are never valid, only arguments are valid, and just as arguments can't be true – because only sentences are

true – sentences can't be valid, because only arguments are valid.

Female: I'm guessing the consistency is all the sentences are true?

Lecturer: Right, it's not that all the sentences *are* true.

Female: Isn't it the consistency then shows us that the argument is bad, or not valid? Something along those lines?

Lecturer: It's certainly the case that when we check the counterexample set for consistency, this will tell us whether the argument is valid or not. But I'm still trying to work out what it is for a set of sentences to be consistent.

Female: All true together?

Lecturer: Is it that they're all true together, or that they *can* be all true together?

Female: Can be.

Lecturer: It's the latter, isn't it? A set of sentences is consistent if there's a logically possible situation in which all are true together, and that situation needn't be the *actual* situation. They needn't actually *be* true. All that's necessary is that they *could* be true

together, and that is a logically possible, situation, where the premises are true, and the conclusion false.

So, if the counterexample set is consistent, if you remember the counterexample set is the premises, plus the negation of the conclusion. If *those sentences* can all be true together, this is a logically possible situation in which the premises are true, and the conclusion is false. Because if the negation of the conclusion is true, the conclusion is false.

You've got to take on board all these definitions, and I realise that they're hard, but they're not hard once you just familiarise yourself with them. Once you learn what they mean, you will be able to play with them, as I do.

There was one other thing slightly wrong: you said with, 'not' put in front. What are the words we put in front of the conclusion to create the negation of the conclusion?

Male: 'It is not the case.'

Lecturer: 'It is not the case that…', well done.

So we evaluate deductive argument as either valid or invalid by taking the counterexample set, which consists of the premises plus the negation of the conclusion, and seeing if they can be consistent. If the counterexample set can be consistent, then the argument is …?

Male: Invalid.

Lecturer: Invalid. No arguments are ever false, invalid, that's right. If the counterexample set is *in*consistent, the argument is?

Female: Valid.

Lecturer: Valid, well done, yes that's right. So, it seems wrong, doesn't it, because consistency ought to go with validity, but I hope you can see why it doesn't? You will see why it doesn't as soon as you understand what 'consistent' means, what 'valid' means, and truth and false and so on.

We also saw that valid arguments can have false premises and false conclusions. Is everyone happy with that? We saw that invalid arguments can have true premises, true conclusions, so that's another rather irritating thing, isn't it?

Just as inconsistency of a counterexample ought to go with invalidity, and doesn't, also you feel like saying that valid arguments ought to have true conclusions, but they don't. Why don't they? What is it for an argument to be valid?

Male: Truth-preserving?

Lecturer: It's for it to be truth-preserving, well done. So, that doesn't mean that the conclusion is true, it means that …? What is it for an argument to be truth-preserving?

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

Lecturer: Good. There is no logically possible situation in which the premises are true and the conclusion false. So if the premises are true, the conclusion *must* be true. That's what it is for an argument to be truth-preserving; that's what it is to be valid.

Of course, even if an argument has false premises *and* a false conclusion, it might *still* be the case that there is no logically possible situation such that IF the premises were true, the conclusion would be also true, or false. I got lost in the middle of that sentence; I hope you can complete it for yourself.

We also looked at how to evaluate deductive arguments from the perspective of everyday life. Because if you remember, all logicians are concerned with is validity. They're really not interested in anything else. What is everyone else interested in?

Male: Relevance.

Female: Soundness.

Lecturer: Relevance, as well, good, and soundness, somebody said. What is it for an argument to be sound?

Female: It must be valid and …

Lecturer: Good, there's an 'and', yes! So, an argument that is sound only if it's valid and if it…?

Male: Has true premises.

Lecturer: Has true premises, that's right. It can't be true itself, of course, because no argument is true.

Good, well done, you did some thinking this week, after last week? Jolly good.

We also looked at the paradoxes of entailment. Can anybody tell me one of the paradoxes of entailment? We had a bit of trouble with these, because they are paradoxes: it's very difficult to understand them. Anyone tell me one of the paradoxes of entailment? No?

Okay, a paradox of entailment would be something like, 'Grass is green, therefore two plus two equals four.' Now, that argument is valid, can anyone tell me why?

Female: Because the statement is true? Both premises are true?

Lecturer: No, because both premises being true is not a sufficient condition for validity, because both premises could be true without the arguments being truth-preserving, couldn't they?

Here, we've got a premise and a conclusion, 'Grass is green' is the premise, 'Two plus two equals four' is the conclusion. Anyone tell me why that argument is valid?

Female: Is it because it doesn't pass the test of 'it is not the case'?

Lecturer: Well it wouldn't pass the test, but no – that's just a definition of what it is for it to be valid, if you like.

[Comment from Marianne: I was wrong to say the paradox of entailment given wouldn't pass the test for validity. Of course it would. Try it for yourself!]

So, let's have another look at it again. 'Grass is green, therefore two plus two equals four', and I've said that that's valid. Which means there is no possible situation in which the premise is true, and the conclusion false. Why does that satisfy that definition?

Female: Because the conclusion is always true?

Lecturer: Because the conclusion *can't* be false, that's right. So, there is no possible situation in which the premise is the true, *and* the conclusion false, because there's no possible situation in which the conclusion is false. Is there? Do you see?

So that, just by definition, falls into the definition of validity, and you might say that that's reasonable. Because surely, a necessary truth should be entailed by everything.

Can anyone tell me the other paradox of entailment?

Male: Tautology?

Lecturer: No that's a tautology, it's a contradiction is used. So, if we've got, 'Two plus two equals five, therefore grass is green', that's the other paradox of entailment. Because if there's no possible situation in which that (the premise) is true, then there is no

possible situation in which that's (the premise) true AND that's (the conclusion) false. So, by default, that has to be valid, which again you might think, 'well anything follows from a contradiction.'

In uttering a contradiction, I don't limit any possibilities, do I? So, if I say, 'I'm wearing a skirt and it's not the case that I'm wearing a skirt', I say nothing, don't I?

There isn't a single possibility that I have ruled out, and therefore I have said nothing, so you might think of that as a way to wrap your mind around that.

Okay, what is it for an argument to be sound? We've done that. Why circular arguments shouldn't persuade us of anything. Why are circular arguments valid? Can anyone tell me? Circular arguments are *always* valid.

Female: Because when the premise is the same as …

Lecturer: Good, that's right. So, 'All whales are mammals, therefore all whales are mammals.' Now, that's valid, because there's no possible situation in which that's (the premise) true and that's (the conclusion) false, is there? How could there be? They're the same thing.

So, they're valid, relevant and, in this case, sound, as well, aren't they? So this should be a really good one. Why might we not notice that this is a circular argument? Can anyone tell me why?

Female: If you add in a lot more premises?

Lecturer: If you add a lot more premises, go on.

Female: I've kind of finished. Are they all saying the same thing?

Lecturer: No, not quite. Can anyone tell me?

Female: It can be distracting?

Lecturer: It would distract from the validity, but go on.

Female: If the conclusion is one of the premises, then it would be valid.

Lecturer: That's right. It doesn't matter how many other things that you add here, because of the monotonicity of validity, it would still be valid, because if the premise is amongst the conclusions, then it doesn't matter how many other premises you add, the conclusion would still be amongst the premises, wouldn't it?

If it can't be the case that that's true (the premise) and that's false (the conclusion), then it *can't* be the case that all the premises are true, and yet the conclusion false. Because one of the premises IS the conclusion, isn't it?

So, we may not notice a circular argument, and we may think, 'Oh that sounds valid, that sounds like a good argument', but we fail to notice that indeed it is valid, but it's valid because it's circular, which isn't a good argument.

Okay good, well done. You've redeemed yourselves.

(Slide 3) Just so see where we are, now, and what we're doing: we know how to recognise arguments now; we know how to analyse arguments now; how to identify their premises and their conclusions; and we know how to set them out logic-book style.

We know how to distinguish deductive arguments from inductive arguments, do you remember? Because you can always tell *a priori* whether deductive arguments are good, or not. Whereas, you need to look at background information to tell when the inductive arguments are good.

Also, last week, we learnt how to evaluate deductive arguments. (Slide 4) So, this week, first we're going to look a bit more about the nature of induction. Then, we're going to reflect on the value of induction, then we're going to look at the different types of induction. Finally, we're going to learn how to evaluate the different types of induction.

So you will have learnt then how to evaluate both deductive and inductive arguments, so you will be able to recognise arguments, analyse arguments, set them out logic-book style and evaluate them.

(Slide 5) Inductive arguments are ampliative, such that they add something that can't be logically deduced from the premises. Do you remember that, with deduction, the conclusion is always included in the premises, because the conclusion is a logical consequence of the premises.

That's why they're either valid, or invalid. Because there can't be any possible situation where the premises are true, and the conclusion false, because the conclusion is *logically entailed* by the premises.

So, if you like, the premises already *include* the conclusion, logically. Whereas, with inductive arguments, that isn't true. We can actually learn something *new* from an inductive argument, so they add something that can't be logically deduced from the premises, which is why we never get certainty from an inductive argument, no matter how good it is.

Inductive arguments are not monotonic. They can always be made weaker, or stronger, by additional information. Do you remember, I looked at - I think this one got you quite exercised - I can't remember what it was, now, but Jones took the knife out of the chap's chest. Or, the policeman saw the knife being plunged into Jones'… Actually, I've got it all wrong, because I can't remember what it was, but do you remember that one?

Male: Yes.

Female: Yes.

Lecturer: Do you remember how whatever I added to it completely changed whether it was a strong argument, into a weak argument, or a weak argument into a strong argument?

That's because with inductive arguments you can always do that. Because none of them are certain; none of them are conclusive.

So whereas with deductive arguments, if they're valid they're conclusively valid. With inductive arguments, they're either weak or strong, and you can change their weakness, or strongness, by adding more information.

Inductive arguments can only ever be evaluated by bringing to bear background information about their subject matter. So, whereas you can look at a deductive argument, and as long as you understand the logical word; the ‘ands’, or the ‘or’s’, or the ‘if’s’, or the ‘only if’s’, etc, you can tell whether the argument is valid, or invalid. You don’t even need to know what the content is.

Do you remember I talked about widgets being something or other? Widgets being what?

Female: Halibers?

Lecturer: Halibers, something like that, yes. It was completely made up word, but you were still able to tell whether the argument was valid, because you don’t need to understand the content of the argument in order to determine its validity, if the argument is deductive.

Whereas, with inductive arguments, you always need to bring to bear your background information, and therefore, you need to understand the content. If you don’t understand the content of the argument, you won’t be able to evaluate it.

Okay, so those are the three properties of inductive arguments, and they all stem from the same thing. (Slide 6) Which is that inductive arguments rely on what this chap here, David Hume, very famous Scottish philosopher, calls the *principle of the uniformity of nature*.

(Slide 7) The principle of the uniformity of nature tells us that the future will always be like the past. So, the future will always be like the past underpins every single inductive argument. If

you think about it, it underpins everything a scientist, for example, will do.

So the core of science is the replication of an experiment. If you do an experiment and you get one result, you will want that experiment to be replicated by other people, and they get the same result, or in biology a similar result, but you want it to be replicated.

Of course, that means you're relying on this premise, aren't you? You're relying on the idea that the future will be like the past? If you do this, under these conditions, you will get this result. Doesn't matter where, doesn't matter when, doesn't matter all sorts of other things. If you do this, under these conditions, you will get that result.

That's the principle of the uniformity of nature. I contrast it with the principle of charity. Do you remember I mentioned the principle of charity to you, earlier?

Female: No.

Lecturer: Didn't I?

Female: No.

Female: Yes, you did.

Lecturer: We can fight it out later, can anyone tell me what the principle of charity is?

Female: Go on, then.

Lecturer: Because I happen to know that you have been at another lecture of mine, where I was talking about charity.

Female: Oh, yes of course you were.

Lecturer: Yes, exactly. Does anyone else think I talked about charity?

Male: No.

Female: No.

Lecturer: No? Right, okay. In that case, I'm not going to start talking about charity now, because I'll mystify you. If we have time, after questions, I'll talk about charity. Actually, no, I'll have to do it now, won't I?

Female: Yes.

Lecturer: We need the principle of uniformity in nature in order to understand the physical world. We believe that the future will be like the past, and that any physical object, if it behaves in a

certain way, under certain conditions, will behave in that same way whenever those conditions are the same.

We don't believe that when we're talking to each other. If I want to understand you, then what I've got to use is what's called the *principle of charity*. Which is the idea that evidence for error is not irregularity, as it is with the principle of the uniformity of nature, it's unintelligibility.

So if I'm talking to you, and you strike me as saying something obviously false, now at that point I can say, 'Well, what an idiot. Not talking to him again', and walk away. Or I can say, 'Oh that's interesting, why are you saying that? Why did you say P? I believe not P. This is *why* I believe not P, why do you believe P?' We might discover that we're both wrong, or I might discover that I'm wrong, or he might discover he's wrong.

We know, with evidence of unintelligibility, that there's evidence for error, but we don't know where that error lies. So whereas evidence for irregularity is evidence for error in the physical world, evidence of unintelligibility, or irrationality is evidence for error in the human world, if you like, in the rational world.

I can say more about that in the questions, if anyone's interested.

What is important for induction is that the principle of the uniformity of nature underpins every single inductive argument.

(Slide 8) Of course, we've got a difficulty, here, because that means that even our strongest inductive arguments, like 'The sun has risen every single day in the history of the world, therefore the sun will rise tomorrow', that's a pretty strong inductive argument. How do we know we're not like Russell's chicken?

Bertrand Russell talked about the chicken who, every day in the whole of its life, the farmer would come out to give it its corn for breakfast, or whatever chickens eat. This day, the farmer came out and the chicken thought, ‘Oh good, breakfast’, and got its neck wrung.

How do we know we’re not in that position with respect to the sun? Answer; we don’t. We have no idea. We expect the future to be like the past, but there’s nothing we can do to argue for the claim that the future is like the past. (Slide 9) Any argument that we attempt to make for it is going to be a circular argument.

(Slide 10) So, if we say that the future has always been like the past in the past, therefore the future will be like the past in the future, okay, you can see how that, again, rests on the circle. It rests on the principle of uniformity of nature, again.

No matter how much we try and justify the PUN, it always comes out in a circle. It can’t be justified, at all.

(Slide 11) Nevertheless, we firmly believe it. Not only do we believe it, Hume argue that we cannot *not* believe that the future will be like the past. It's actually impossible for us not to expect that this will happen again, if in the past that has always happened.

Hume thinks that it's part of the very structure of human thought, but, he says, ‘It's not rational.’ By that, he didn’t mean that it's *irrational*, he meant that it's *non-rational.* So, the opposite of rational is not really irrational, because in order to be irrational, you’ve got to be rational haven’t you? This table can’t be irrational, can it?

Female: No.

Lecturer: That's because this table is non-rational, it just doesn't engage in reasoning. I engage in reasoning, therefore I can be irrational. So, irrationality is a failure, if you like, in the house of reason, and non-rationality is the opposite of rationality.

Hume thinks that inductive reasoning is not, in a way, reasoning, at all. It's not rational, it's just something that we are forced to do, because that's the way our brain is structured. Very interesting, isn't it?

(Slide 12) Some philosophers have believed that this shows we should reject induction as a means of acquiring knowledge. After all, surely knowledge should be based on reason? Reason is what leads us to knowledge, and if inductive argument isn't rational, if it's not rational at all, then it shouldn't be. It's not a way of acquiring knowledge.

(Slide 13) Karl Popper, who's a very famous philosopher of science, he argued that the only good arguments are deductive arguments. (Slide 14) He said that a scientist should never try and confirm hypotheses, and they should never claim to know that a theory is true. Instead, they should try and *falsify* hypotheses, and at most, they should say of a theory, is that it hasn't yet been falsified. So they should never actually say it's true.

The idea behind this is that when you try and confirm something, you're using deductive arguments. So you're saying, 'If all swans are white, then the next swan I see is going to be white.' If the next swan you see is white, then there's a little tick (against your hypothesis).

There are all sorts of problems with the logic of confirmation, and somebody wants me to talk about grue. Again, if we have

time in questions, I'll talk a bit about grue, and the rest of you can just wonder what on earth I'm talking about.

Let me instead quickly give you the ravens paradox; if all ravens are black, that should be confirmed by the sighting of a black raven, shouldn't it?

So, the sighting of a black raven gives you in a tick in the box beside 'all ravens are black'. Would you think that 'all ravens are black' is also confirmed by the sighting of a white gym shoe? Would you?

Female: Not immediately, no.

Lecturer: Not immediately, okay, but actually 'all ravens are black' is logically equivalent to 'all non-black things are non-ravens', isn't it? 'All ravens are black' is logically equivalent to … let me write that down…

Female: Doesn't say exclusively ravens are black though, does it?

Lecturer: No, let me write it down so you can see what I'm saying. Can you all see this?

Male: Yes.

Lecturer: 'All ravens are black' is logically equivalent to 'all non black things are non ravens'.

Male: Yes.

Lecturer: Do you accept that?

Female: Yes.

Lecturer: ‘All non black things are non ravens’. So what the logical equivalence means is that (‘all ravens are black’) couldn’t be true unless that (‘all non-black things are non-ravens’) is true, and that (‘all non-black things are non-ravens’) couldn’t be true without that (‘all ravens are black’) being true. These are true and false in exactly the same conditions. Fair enough?

Male: We should mention, you have things in the first sentence. Should it not be ‘all ravens are black things’, and therefore ‘all non-black things are non-ravens’?

Lecturer: I can put it in if you like? I don’t think it's really necessary, because you just mean they’re coloured black, don’t you? Here, you’re saying all things that aren’t black…

Male: We’ve had to identify what are non-black in the second sentence.

Lecturer: That's because here I'm talking about all ravens, whereas here I'm talking about non-ravens, aren't I? So, I need the things in there. Anyway, I think that's a bit of a red herring.

Given that a white gym shoe is a non-black, non-raven, isn't it? Well, if 'all ravens are black' is confirmed by the sighting of a black raven, it ought also to be confirmed by the sighting of a white gym shoe, which is a non-black, non-raven. Do you see the problem?

Male: I just hate that.

Lecturer: You hate it? Why do you hate it?

Male: It's almost like this thing where all Chinamen wear brown shoes and if you wear brown shoes you're a Chinaman, no?

Lecturer: Well, it would have to be all *and only* Chinamen wear brown shoes. Therefore, if you wear brown shoes you're a Chinaman.

Male: I can see how it's absolutely fine, the non-black thing, so it's a non-raven, but the sighting of a white gym shoe does not confirm that ravens are black.

Lecturer: Okay, why not?

Female: It's not very persuasive, is it?

Lecturer: Why not? You're absolutely right, none of us think for one minute that the sighting of a white gym shoe is any confirmation of the theory that all ravens are black, but the logic says it is.

As far as logic is concerned, the sighting of a white gym shoe should be every bit as much a confirmation of 'all ravens are black', as the sighting of a black raven.

So what is it that means that it's not persuasive? Because you're quite right, it's not persuasive.

Female: I would say it's because of probability.

Lecturer: Go on.

Female: The number of things that are not ravens, I would say in my mind, is far more than ravens.

Lecturer: Good, well done. What you're doing is you're bringing to bear your background knowledge of the world to tell you that actually this is a world in which there are many, many more non-black non-ravens, than there are black ravens.

Imagine a world in which there are just 11 things. You know that this is a world where there are only 11 things, and I know there are only 10 things, and 9 of them are ravens, and 9 of them are black. In that world, the sighting of a white gym shoe would be *conclusive* confirmation of the claim that all ravens are black, wouldn't it? Do you see?

The Ravens Paradox it's called, and it was a chap called Carl Hempel who came up with it, what that shows you is the necessity of background information. We unthinkingly bring to bear our background beliefs about the world in deciding what does test our hypotheses and what doesn't.

So, you have a hypothesis: 'all ravens are black', you think, 'I know what confirms that, it's the sighting of a black raven', and what you don't realise is, you're smuggling in background beliefs about the world.

Actually, I am going to do 'grue', I may regret this, but let's have a go.

Female: Marianne, sorry, just before you leave that one, all that would make sense to me; the first sentence said, 'All black things are ravens.'

Lecturer: No, it doesn't.

Female: No, I know it doesn't, therefore to me, it doesn't make sense. Because by saying 'all ravens are black', ergo there was no such thing as a not black raven. All non-black things, any other colour, is not a raven. I still think that allows room in the world for things that are black, and are not ravens. It doesn't exclude it, those two sentences.

Lecturer: For things that are black and not ravens?

Female: Yes, they're black but aren't ravens.

Lecturer: Those are irrelevant to this, aren't they?

Female: It isn't when you bring in the gym shoe.

Lecturer: No, because the gym shoe is a non-black, non-raven. It's not a black non-raven.

Female: No, it says 'all ravens are black', it doesn't say 'every black item in the world is a raven'.

Lecturer: What you're saying is it doesn't say there are… So, this is consistent with there are black non-ravens?

Female: Yes.

Lecturer: I.e., Marianne's boots?

Female: Yes.

Lecturer: Okay, is that what you're saying?

Female: Yes.

Lecturer: That's true, but why is that relevant to this? Because this is saying 'all *ravens* are black'?

Female: Yes.

Lecturer: Can I leave out the 'things', it bothers me? Okay, 'all ravens are black', so 'all non-black things are non-ravens'. It doesn't say all black things must be ravens, at any point does it?

Female: No.

Lecturer: You're getting your Venn diagrams mixed up.

Female: I thought that. So, you see a white gym shoe, you know it's not a raven ...?

Lecturer: And you know it's not black.

Female: You definitely know it's not black.

Lecturer: My black boots are completely irrelevant to the confirmation either of 'all ravens are black', because my boots aren't ravens. It's also irrelevant to the confirmation of 'all non-black things are non-ravens', because they're not non-black.

Female: Right, I'll digest that for a while.

Lecturer: Does anyone else agree with me? Does anyone want to argue with that? What you're doing is you're making the mistake of … No, I'm not going to into that.

Female: Don't worry …

Lecturer: Yes, you'll discuss, thank you very much.

I'm just going to quickly do grue.

Male: Sorry, Marianne, I have a question. Just before, when you said there are only 10 things, and 9 of them are all black, and one of them wasn't…

Lecturer: There are only 10 things, 9 of them are black and 9 of them are ravens.

Male: That seemed to be, when you talked about the Venn diagram there, effectively you have a set that you know, and you have one thing that's in a set. So you know that that set is not a sub set of that, they stand alone. Well, that's only because you know the population of the set.

Lecturer: Well, there are 10 things, okay? Nine things are ravens, so there is one non raven?

Male: Yes.

Lecturer: Nine things are black, so there is one non-black thing, okay, but these needn't be the same thing, need they?

Male: No.

Lecturer: So you could have a white raven in this world, but if you see a white gym shoe ...?

Male: You know by exclusion that there must be nine black ravens.

Lecturer: Exactly, you know that *all ravens are black*. Do you see?

Male: Yes.

Lecturer: So the idea is that when you bring to bear your background knowledge of that world [the world with ten things], you know that seeing the white gym shoe confirms 'all ravens are black'. Whereas, when you bring to bear what you believe about *this* world, you know that single white gym shoe ***doesn't*** confirm all ravens are black.

The moral of the tale is we don't even notice that we're bringing to bear that background information, when we decide what will confirm our theories.

Here’s another paradox of confirmation; all observed emeralds, every emerald we’ve ever observed, has been green. I don’t know if that’s true, or not, but it's irrelevant whether it's true, or not. Let’s just say all observed emeralds have been green.

Therefore, we going to form a hypothesis ‘all emeralds are green’. A chap called Nelson Goodman, wanted us to say, ‘Okay, well if we define ‘x is grue’ as ‘x is green and observed, or blue and unobserved’, then we have as much reason to think that ‘all emeralds are grue’, as we do for thinking ‘all emeralds are green’.

You might think, ‘Hang on a second, ‘grue’ is a made up predicate.’ Well, yes okay, lobsters, we know, are ‘blink’ aren’t they? They’re either brown and uncooked, or pink and cooked, aren’t they? So, lobsters are blink, and there’s nothing wrong with blink, so why should we think there’s anything wrong with grue?

Do you see that if you define grue thus, we do have every bit as much evidence for the truth of ‘all emeralds are grue’, as we do for ‘all emeralds are green’. Because every emerald we have ever seen has been grue, hasn’t it?

What’s more, every emerald we see from now in is going to confirm ‘all emeralds are grue’, isn’t it? That’s because the minute we see it, it will be green and therefore grue. Do you see what I mean?

Exactly the same with the ravens thing. We absolutely know, or we take ourselves to know, that emeralds are not grue. We just won’t project the predicate ‘grue’, and yet we will project the predicate ‘green’. Well, why?

Why do we think that ‘green’ is an acceptable predicate, but ‘grue’ isn’t? There’s no doubt that we do think that, but why do

we think that? Because actually, again, the logic tells us there's no difference between 'all emeralds are green', and 'all emeralds are grue'.

Male: Yes, that makes it a lot clearer, because every time I come across this argument, locked into the argument, there seemed to be date involved. Before a particular date, and after a particular date, and that really confused me, because it seemed too arbitrary in many ways.

Lecturer: Yes, sometimes people put a date in here and say 'x is grue, if x is green before t, or blue after t', and actually that just confuses everything. This is much better.

Male: Yes, certainly did.

Lecturer: Can anybody tell me why we don't project 'grue'? You can tell, I hope, that the answer is going to be comparable to the ravens puzzle.

Female: Well I would give a different answer; I would say that 'grue' or 'green' are just labels, and we know there is a content for 'green', and we know there's no content for 'grue'.

Lecturer: No, I dealt with that with the 'blink'. No, you don't know there's no content to grue, you don't know that. We have as much evidence for 'emeralds are grue' as we do for 'emeralds are green'.

Female: It's based on the preponderance of the evidence the word ‘grue’ doesn’t exist.

Lecturer: I’m sorry, the word ‘blink’ is a perfectly good predicate, isn’t it? There is nothing wrong with the predicate ‘grue’. That’s not the answer.

Let me tell you the answer, because otherwise it will go on too long.

Male: Is it to do with the belief in the uniformity of nature?

Lecturer: All of this is to do with the belief in the uniformity of nature (the PUN), but what I’m asking, I suppose, is which bit of nature is deemed to be uniform in this case?

The answer is that we have no reason, do we, to think that things change colour when we observe them? Or things like *emeralds* change colour when we observe them. We do have reason to think that lobsters change colour when we cook them.

So we’re quite prepared to accept that lobsters are blink, but we don’t have any background beliefs about emeralds changing colour when we look at them. So, although we have every bit as much evidence for ‘emeralds are grue’ as ‘emeralds are green’, our background knowledges come to bear, and says, ‘No, actually, all emeralds are green’ that’s the one we project and form as a hypothesis, not ‘all emeralds are grue’.

I think you preferred the ravens one, didn't you? Maybe I shouldn't have tried grue.

Popper, who is completely different, and has - well, when I say nothing to do with this, of course he did, he was a philosopher of science and a very good one - but he was saying something different. What he was saying is we should never try and confirm hypotheses. All these paradoxes are reason to think that actually confirmation is just not the way we're going. One tick every time you see a black raven, or a white gym shoe.

Instead, scientists should try and *falsify* hypotheses, and they say of theories that they haven't yet been falsified.

(Slide 15) Actually, in saying this, Popper illustrates Hume's claim that we can't manage without induction. Why should we believe, after all, that a hypothesis that has been falsified at t will again be falsified at t+1? We can't reject a theory just because it has been falsified, because that assumes that we think it will be falsified next time we test it. Why would we think that, if we weren't using induction, which is what we've been told we shouldn't use.

So, we're relying on the PUN, even when we're trying to falsify things.

(Slide 16) Popper was wrong to think we could give up induction. Anyone, and that may include the people here, who think that induction is somehow inferior to deduction, because it doesn't give us certainty, you're manifesting your failure to understand that without induction, we couldn't learn anything *new* about the world, at all.

Induction isn't inferior, I can't say that often enough, it's just *different*. So there's deduction and induction; deduction gives us certainty, induction never gives us certainty, but that doesn't mean that induction isn't a way of acquiring knowledge. In fact,

it's our *only* way of acquiring knowledge about our actual world, about the world in which we live.

Male: You've just said exactly that point, because you made the point that induction can never be certain, whereas deduction is always certain, well validly.

Lecturer: It's either certain or of no use at all.

Male: Yes, exactly, I think that was Popper's argument.

Male: How does the level of certainty come into this? Because there may be situations where you can actually say, 'There's a probability of that outcome being generated. We've identified 50% probability, but we haven't identified another 50%.'

Lecturer: I'm sorry, I don't understand the question.

Male: What is the level of certainty?

Lecturer: Deduction will give us absolute logical certainty. We know beyond all possible logical doubt that if a deductive argument is valid and the premises are true, the conclusion *will* be true. There is absolutely no doubt about that, it will never change, it's completely conclusive. We've got absolute certainty about it.

Whereas, the truth of an inductive conclusion can never be established, certainly on the basis of an inductive argument. Never. It doesn't matter how good the argument is.

Male: Can there ever be a level of certainty?

Lecturer: There can be a different level of probability. If you remember going back to when we were looking at the difference between induction and deduction; deduction is an either/or matter. It's either valid, or it's invalid, end of story. Whereas, induction is always a matter of degree. It's either weak or strong, and therefore everything in between, it can be any of those things.

Of course, background information can be brought to bear to change that probability level.

(Slide 17) So the three qualities of induction mentioned above, or before - I don't think that's really a word – 'ampliactivity', but never mind, non-monotonicity (and that's certainly not a word), a-priocicity - so far they've made it impossible for us to systematise the evaluation of inductive arguments in the way we've systematised the evaluation of deduction.

What I gave you last week was one method for testing any deductive arguments. You can use that same method and it will always show you whether a deductive argument is valid, or invalid. We can't do that with induction.

(Slide 18) Deductive validity is extremely well understood. It's annoying that it throws up the things like the paradoxes of entailment, but actually they don't bother us, and they don't stop us creating computers that get better and better and better, all the time.

Inductive strength isn't very well understood, at all. Notwithstanding the mathematical theory of statistical probability, there are very good accounts of probability, that are systemised in the way that deduction is systematised. I'm not going to talk about that, because we're not doing mathematics, here, and I don't know enough about it.

(Slide 19) What is clear, though, is that there are *strong* inductive arguments and *weak* inductive arguments, and everything in between. We know a lot about how to evaluate the different types of induction, and that's what we're going to do.

(Slide 20) Before we look at the different types of induction, we should remind ourselves that we always want the premises of our arguments to be true. (Slide 21) So, just with respect to deductive arguments, we always ask these two questions; 'Are the premises true?' 'Is the argument valid?'

With inductive arguments, we always ask, 'Are the premises true, and is the argument strong?' So, those are the two key things. When you come to evaluating arguments, those are the two key questions that you ask yourself.

(Slide 22) Notice I've said, 'I classify inductive arguments thus'. If you go onto the web, you will see that there are lots of different classifications and I just gave up trying. I thought, 'I'm just going to decide how to do it myself.' That's how I do it.

Some people think that abductive arguments are neither inductive or deductive, I don't agree with that. I think they're inductive. Again, if you go on the web, you'll see, and I've given you some references at the end of this week, which will help you on that.

(Slide 23) That's how I classify them, and what we're going to do is we'll look at each type of argument, making sure we

know how to recognise it, and then we'll go back to each one and learn how to evaluate it.

(Slide 24) First of all, *inductive generalisations*; that is the first type of inductive argument. This type of argument makes claims about whole populations from observations of samples of those populations. So, every swan I've ever seen has been white, therefore all swans are white. Or every raven I've ever seen has been black, therefore all…

Female: Gym shoes are white.

Lecturer: All non-black things, are non-white. Do you see how that's an inductive generalisation? It takes some observations and extrapolates the observations to the rest of the population, or the rest of the things that have been observed.

I realised last time I did these lectures that lots of people get confused about this. (Slide 25) Deductive arguments have premises that must be based on induction. So, this argument is a deductive argument; 'all swans are white, the bird in the next room is a swan, therefore the bird in the next room is white'. Do you see that that's deductively valid?

That premise (premise one) must be an inductive premise, mustn't it? The only way we could have got to that is on the basis of an inductive argument.

What this shows you is that actually science is based on both induction and deduction, because you need your induction to get your hypotheses, but then you need your deduction to get the test that you're going to carry out on your premises. So don't ever think that induction is what scientists do, they don't

go for deduction. Of course they do, they have to for this reason.

That’s a deductive argument, even though the first premises is a classical claim from induction.

(Slide 26) Okay, this is the second type of inductive argument; *causal generalisations*. These makes claims about causation, on the basis of claims about correlations. This is just a type of inductive generalisation, because it's taking observations and extrapolating them into the future.

So, what we do is we have seen A and B always correlated, therefore they will always be correlated in the future. We call ‘always correlated’ in that way, causation, or we often do. Our evidence for causation is constant conjunction, or a correlation.

Here, we might say, ‘Many of the people at Jane’s party on Saturday were ill on Sunday, their illness must have been caused by something that they ate at Jane’s place.’ That’s a causal generalisation. I probably shouldn’t have put the word, ‘must’ in there, because it's going to confuse you and make you think it's deductive, but it isn’t.

So, we’ve got a correlation here, and we’re saying that that’s evidence for cause, that’s a causal generalisation.

(Slide 27) *Analogies* are also inductive arguments. They make claims about things on the basis of claims about their similarity to other things. ‘So the universe is like a watch, watches have makers’ - we’ve never seen a watch without a maker – ‘therefore, the universe must have a maker’. Do you see? So, it's extrapolating the similarity, or rather on the basis of the similarity.

(Slide 28) Arguments from *authority* make claims about the correctness of things someone says on the basis of that person's being an authority. So, if you like they've said things that have been correct in the past, therefore in the future, they're going to say things that are correct. Do you see how that rests on the principle of the uniformity of nature?

Einstein says that, 'Nothing goes faster than the speed of light in a vacuum, therefore nothing goes faster than the speed of light in a vacuum'. 'Einstein says P, therefore P is true'.

(Slide 29) *Abductive* arguments, finally; they're arguments to the best explanation, you've probably heard them called before. Abductive arguments; they make claims about which explanations are likely to be correct on the basis of observations about the sorts of explanations that have been correct in the past.

So, you might say, 'Crop circles are eventually going to be explained by science, because in the past, many things thought to be susceptible only to supernatural explanations have been explained by science.' Here, you're saying that, 'Scientific explanations are good explanations, and they will eventually explain crop circles. '

Male: Can we go back to the watch-maker, please?

Lecturer: No, because I'm just about to get on to evaluating these arguments. I haven't looked at how to evaluate them, yet. All we've done is learned how to recognise them.

Male: I wanted to argue that is a deductive argument.

Lecturer: That's a deductive argument? Why?

Male: Premise one, 'the universe is like a watch'. Premise two, 'watches have makers', conclusion 'therefore the universe has a maker'.

Lecturer: Is there any possible situation where that's true (premise one), that's true (premise two) and that's (conclusion) false? Can anyone think of a situation which that's true, that's true, and that's false?

Male: Yes.

Male: Yes.

Lecturer: Yes? Okay, does that mean that that's not a good argument? Does that make that argument invalid? No.

Male: No.

Lecturer: No. Well, it's not a deductive argument, it's an *inductive* argument. Do you see why?

Male: Because it's *like* a watch, not is a watch?

Lecturer: Yes, exactly. If it said, ‘The universe *is* a watch, watches have makers therefore the universe has a maker’, that would be a deductive argument. Yes.

Male: It's because it's an *analogy*.

Lecturer: Yes, well…. Right, okay. Have we done that?

All: Yes.

Lecturer: We’ve done that?

Female: Yes.

Lecturer: (Slide 30) Can we classify these arguments by type? ‘Many undergraduates like to get drunk at the weekend, so I should imagine that as it's Saturday, James may be in the pub.’

Male: Generalisation.

Lecturer: It's a generalisation, an inductive generalisation, yes good. What about, ‘Jessamy is the Chief Executive Officer, so if she said expenses won’t be paid today, then expenses won’t be paid today’.

Female: Authority.

Lecturer: Authority, yes. ‘Often, after I’ve been in an aeroplane, I’ve developed a cold. I’m travelling to Johannesburg on Saturday, so I’m expecting to have a cold by Tuesday.’

Male: Causal.

Female: Causation.

Male: Causation.

Lecturer: Yes, I wasn’t sure about that one myself. It's either inductive, or it's causal, isn’t it? Either way, it's one of the two.

‘Bankers are like vampire bats; they suck our blood and should be destroyed.’ I apologise to any bankers.

Male: Just true.

Lecturer: Just true won’t do! What sort of argument is it?

Female: Analogy.

Lecturer: It's another analogy, that's right.

'The last three times I've asked Susan what was on at the cinema, she misled me, so I shan't ask her again.' There's an argument implied there isn't there? Rather than an argument.

Male: Causal generalisation.

Female: Inductive generalisation.

Lecturer: It's an inductive generalisation. Yes, I'm not sure if causal is…

Male: I'm just correlating. If I'm correlating from previous accounts.

Lecturer: Yes, correlations are taken as evidence for cause. Husbands and wives are correlated, they don't *cause* each other.

'Many children have developed autism shortly after having been given the MMR jab, therefore the MMR jab causes autism.'

Female: Causal.

Lecturer: That's a causal generalisation, yes. Unfortunately, that was the one that was underneath Andrew Wakeman's claim, which led to the loss of herd immunity. Because many children have developed autism shortly after having been given the MMR jab, but that's because the MMR jab is given at roughly the

time children develop autism. Or just before the time that children develop autism. There is a correlation there, but no causation.

(Slide 31) Next, we're going to learn how to evaluate those arguments. We're going to learn this by listing the questions to which we'd like answers and see why we'd like answers to them.

After each one, I'm going to include an exercise, but this is this week's homework exercises. So you will get the answers to them next week, we won't do them now, which is just as well.

(Slide 32) We'll always be asking, assuming the premises of the argument are true, how strong a reason are they to believe that the conclusion of the argument is true?

So, if you remember the two questions I gave you at the beginning, 'Are the premises true?' 'Is the argument strong?' this is the second part, that second question that you should be asking.

(Slide 33) Evaluating inductive generalisations; 'every swan we have ever seen has been white, therefore every swan is white'. Firstly, you need to recognise what's the sample, and what's the population.

What is the sample, here? It's the swans we have seen, isn't it? So all you could say is that, 'Every swan I've ever seen has been white. I assume that every swan I've ever seen is representative of all swans, therefore all swans are white.' So that's the sample, and that's the population.

You need to know something about how large the sample is. How many swans have we actually seen? Who is this, 'We'? Is it *me and my pal*, or is it *all English people*, for example? Or *all scientists*, or *all naturalists*? Who is it?

Have we any idea how the number of swans we've seen compares with the number of swans there are? The number of the population taken as a whole.

The next question we ask is, 'Is the sample representative of the population?' Do we think there's something odd, or special, about the swans that we've seen?

A very popular example was used in an American election, in which they did a poll, 'Who is going to vote for..?' whoever it was. When it came back, is said something, 'Sixty percent are going to vote for this person.' In the result, it was something like 10%, and what had turned out is that they had conducted the poll by *telephone*.

Actually, very few people owned telephones in those days, and the people who owned telephones tended to be the rich people, who were going to say, 'Yes, they would vote for this person'. Whereas those that didn't have phones, the vast majority, were less likely to vote for that person.

So, you need to know whether your sample is representative of the population, because if it isn't, then what you're going to get out is - well you've heard the saying - 'Rubbish in, rubbish out.'

Also, you want to know if there are any counterexamples. This, 'We' here, who is it? *We* have seen only ever white swans, but has anyone else seen any other type of swan?

(Slide 34) So, the claim, 'all swans are white' was seen to be false in 1790, when John Leyton sailed up the Swan River, it's called - for that very reason - in Western Australia and discovered black swans. Do you think that his discovery meant that the argument for the claim that all swans are white was a bad argument?

Female: No.

Male: No.

Lecturer: It's interesting isn't it? Because we now know that the conclusion of that argument is categorically false. It's not the case that all swans are white, but actually it wasn't obviously a bad argument, was it, in the first place? Because at the time, we had seen many swans. We had seen them in lots of different places, lots of people had seen different ones. All these answers could have led us to think that actually this was quite a good argument.

Female: So it was an example of Popper's principle being useful, up until that point?

Lecturer: Yes, and in fact a book has just been written called Black Swan Events; Black Swan events are the falsification of something that everyone takes to be absolutely normal.

Do you see how induction works? All the evidence you have in suggests that this hypothesis, 'all swans are white', is true, and then one falsification of it comes up and shows that actually, it (all this evidence) doesn't matter. So, a bit more information and the whole argument is just blown.

So, actually it doesn't necessarily mean that it was a bad argument.

Any questions about inductive generalisations and how to evaluate them?

Female: It's not statistical; you only need the one.

Lecturer: It's certainly the case with a counterexample. Just as you have counterexamples to deductive arguments, once you've got the word 'all' in there, you can come up with a counterexample.

If we're concluding *many* swans are white, the sighting of a black swan wouldn't have falsified it.

Female: David Deutsch reckons that the sorts of arguments or debates you're talking about, particularly extrapolation, is very post enlightenment. It's based on empiricism isn't it? 'We have seen that to be true, therefore it is true.' He argues that that is only, that's the paradigm we work to, because we're all post-enlightenment and that it wasn't true before, and it shouldn't be in the future. Discuss.

Lecturer: Well, I don't want to comment on David Deutsch, because I haven't heard him make this argument, but I do want to say it's certainly the case that in the enlightenment, we started to take note of observation, and things.

For example, up until the enlightenment, people were saying, 'Aristotle said, 'P', therefore P.' End of story. So an argument from analogy was deemed perfectly acceptable, as an argument.

There's a wonderful example of, I can't remember his name now, but it was argued that the number of planets was seven, because there were seven ways of knowing about the world. I can't remember exactly how it went, and I might have used it later on.

During the enlightenment we became very interested in empirical evidence for things, i.e., correlations. So the principle of the uniformity of nature became very important, but don't think for one minute that inductive argument can take over from deduction. Deduction is every bit as important as induction, and they're quite different.

Actually, this is quite a good question, because I can get quite annoyed when I see people in the newspapers talking about, 'We need evidence for this'. Actually, I think it should be, 'We need evidence or arguments.' Because, there are certain things for which you will never have evidence, because you argue for them. Deduction is the only way to show that they're true.

So, we have evidence *and* arguments, it's not *just* evidence, so deduction is never going to go away. That's just nonsense.

Male: What do we do about if the population sample size is none, but what if we don't understand or know how big the population is?

Lecturer: We should in that case probably hedge our bets on the conclusions we make. So, if we had said, 'many swans are white', maybe even, 'most swans are white', we wouldn't have had our conclusions falsified by the sighting of a black swan, would we?

So, we tell students that they shouldn't make conclusions stronger than the argument can bear, and that's true here, as well.

Do you remember I talked about inductive boldness at some point? Different people are differently inductively bold. So you (pointing at one student) might come in and see me wearing earrings once, and think, 'Oh, Marianne always wears earrings'. Whereas, you'll (pointing at another student) want to see me 100 times before you think, 'Oh yes, she always wears earrings.' So you're more inductively bold than you are. You need more evidence for your claims.

Female: I just wanted to ask, in this current argument we are moving from logical possibility to physical possibility, with the whole business of swans?

Lecturer: Well, you could see it like that. Whether there's any such thing as physical necessity is a big question. We know there's logical necessity. David Hume had things to say about physical necessity.

So, our only evidence for causation is correlation. We see two things that are always correlated, one is always prior to the other, etc. Actually, I'm going to talk about this a bit when I get into causal generalisation, so would you mind if I just carried on?

Female: Okay.

Lecturer: If I don't answer your question, come back to me again. Two more questions, then I had better move on.

Male: In regards to the swans one, in practice, what difference does it make?

Lecturer: What difference does what make?

Male: Well, I'm just about to ask you. 'All swans are white' was seen for hundreds of years to be upheld by the experience of mankind. Then, in a very particular circumstance, there was an exception. At that point, it changed from 'all swans are white' to 'most swans are white', but there was another statement that replaced it.

Lecturer: Yes, are all swans in the southern hemisphere..?

Male: Yes, that is likely to happen with all inductive arguments. You make a claim, and that claim will carry on being held as the experience of truth for the population, until it bumps up against another reality. At that point, it changes.

Male: It's degrees of probability.

Male: Yes, saying this was false…

Lecturer: Well, it's only false because it's '*all* swans are white'.

Male: 'All swans are white' was true.

Lecturer: No, it wasn't. Hang on, it was no more the case 'all swans are white was true', than 'that the earth was flat' was true when everyone *believed* the earth was flat. We believed it was true, but it *wasn't* true.

Male: No, I'm not following you down that path.

Lecturer: You're not following me down that path?

Male: 'All swans are white'; if I had a cage that has got swans in it, and all the swans in it are white, and I say 'all swans are white', it's true then, isn't it?

Male: No.

Female: No.

Lecturer: It might be 'all the swans *in my cage* are white', might be true, but not '*all* swans are white'.

Female: Or the swans in the world.

Male: Yes, but if the experience of human nature, up until 1790…

Lecturer: He's playing right into my hands, I'm enjoying this. Go on?

Male: Alright, if up until 1790 the experience of human life, or animal life, was that all swans are white, it was true. There was no reason to believe otherwise.

Lecturer: What's your name?

Male: I'm not going to give you it.

Lecturer: Stop laughing! What's his name? Go on, what's your name?

Male: 'William'.

Lecturer: Okay, you've given me the opportunity to do something I love doing, so I think this is wonderful.

'William believes Marianne is wearing red', is that true?

Male: Today.

Lecturer: Yes. God, you are nitpickers, aren't you?

Male: The whole theory is about nitpicking.

Lecturer: Yes, I know, you're nitpickers amongst nitpickers. Okay good, is that true?

Male: Yes.

Lecturer: Good, okay. I want you to notice that there are two sentences here; there's one big sentence and it embeds another sentence, here. Sorry, it's a shame it's such a long sentence, but the one sentence, 'William believes Marianne's wearing red', and embedded in that, there's another sentence, 'Marianne's wearing red' Is that fair enough?

Female: Yes.

Lecturer: Now, could that one 'William believes Marianne's wearing red',be true and that one 'Marianne's wearing red' be false?

Male: Yes.

Lecturer: Yes, could that one 'William believes Marianne's wearing red',be false and that one 'Marianne's wearing red' be true?

Male: Yes.

Lecturer: Yes, so Marianne's wearing red, but William doesn't believe it. He didn't come today, so he's not seen me, so he doesn't have any beliefs about what I'm wearing.

Could they both be true?

All: Yes.

Lecturer: Could they both be false?

All: Yes.

Lecturer: What that tells us is that the truth value of that belief ('Marianne's wearing red') differs quite independently from the truth value of the embedding belief.

The only way you'll miss that is if you take the expression, 'Marianne's wearing red is true for William.' Now, that's ambiguous. One might be just that, 'William believes Marianne is wearing red'. 'Marianne's wearing red is true for William' just means 'William believes', therefore '*William believes it's true that Marianne's wearing red'*.

The other one, the sentence, 'Marianne's wearing red is true for William even if not for anyone else.'

Now, that (the first one) is just straightforwardly acceptable, isn't it?

Female: Yes.

Lecturer: We understand that, we know exactly that people sometimes form beliefs for all sorts of reasons.

That one (the second) is a very strange thing to say, isn't it? If we say, 'Marianne's wearing red is true for William, even if not for anyone else'. In other words, not that it's just that he believes it, but that it's *true* for him. Surely, either I'm wearing red, or I'm not and if I am wearing red then it's true for *everyone*, not just William. If I'm not wearing red, then it's not true for William either, even if he believes that I am.

Do you see where I'm getting to, William? What I'm saying is that when everyone believed that the earth was flat, that didn't make it *true* the earth is flat. What it meant was that everyone *believed* the earth was flat. Do you see? That they *believed* it was true that the earth was flat, didn't make it *true* that the earth was flat.

This is an absolutely crucial logical point, that actually many, many people miss, and makes a huge problem for thinking about epistemology, metaphysics, anything to do about truth, in fact. Have I convinced you?

Male: I accept that entirely, but the point I was making was in practice what difference does it make? Because you carry on believing it's flat, until it's shown to be round, and then you move on to a new truth basis.

Lecturer: Well, no you move onto a new belief at that point.

Male: A new belief?

Lecturer: Yes, exactly. So, what difference does it make? Well actually, because everything we do is a result of what we believe and because, if our beliefs are true our actions are usually successful - in other words, the truth of our beliefs, we like to think leads to the success of our actions, the falsehood of our beliefs leads to unsuccessful actions - we care whether our beliefs are true. So it does make a difference.

All that has happened here is that we over-generalised, but actually we could probably say that we had every good reason to over-generalise. In fact, that human beings are given to over-generalising.

We're generally much more inductively bold than we should be, and one of the things about being a scientist is that you should rein in your inductive boldness. You shouldn't say that you've shown something, when in fact you haven't.

Male: Okay.

Female: In defence of William, surely all inductive arguments are open to what you've just said, and it's relevant?

Lecturer: Yes, they are. Inductive arguments never give us certainty, that's right, but we got to the moon on them. There were quite a few inductive arguments involved in the calculations that took us to the moon. Science is based largely on induction.

Just because it doesn't give us certainty, doesn't mean it doesn't give us knowledge.

Female: It gives us a starting point, and then we get curious.

Lecturer: It probably gives us more than a starting point, as well. It takes us to somewhere we otherwise couldn't have gone.

Can I move on? Goodness, we've only done inductive generalisations. (Slide 35) There are some to do at home.

(Slide 36) Okay, the evaluation of causal generalisations, now. 'Many of the people at Jane's party on Saturday night were ill on Sunday. The illness must have been caused by something they ate at Jane's party.'

(Slide 37) Questions we should ask; firstly how many correlations have been observed? It says, 'many' here, 'many of the people at Jane's party…' implying that not all of them. So there were some people at Jane's party didn't get ill. So, how many correlations were there between people who were at Jane's party and people who got ill?

Under what circumstances have correlations been observed? Was it Freya, who wasn't invited to the party and doesn't like Jane, who tells us that a lot of the people who were at Jane's party got ill? If so, we might question that.

I suppose, we're assuming that the premises are correct, aren't we? So maybe that's unfair, but you can see we would ask that.

Are there any exceptions to the correlation? Are there any who are ill who weren't at Jane's party? We know that there were some at Jane's party who weren't ill, but are there any who

were ill who weren't at Jane's party, we might ask? That's continued.

Could the correlations be accidental? Could the illnesses have had variable causes? So it was just an accident that all the ill people were at Jane's party, or mostly ill people were at Jane's party.

Could the correlations be explained by common cause? Did Jane throw her party because of an earlier event at which those were ill might have caught something, or even something bad? So, maybe Jane had the party because we all went to the Olympics that day, and we ate a hamburger and it made us all ill. Jane's party was actually nothing to do with it.

Could the causal relation run the other way? Could it be that it was because they were going to get ill that certain people went to Jane's party? I'm sure there must be people here who have run chicken pox parties, is that right? So, the people who went to that party went to the party because they *wanted* to get ill, didn't they?

Does the causal relation make sense? Pineapples and striking matches; this is how it was taught me and I continue to use it numerous years later. If it had been, in the history of the universe, true that every time a pineapple dropped from a tree a match was struck, would we think that there was a causal relation between the dropping of the pineapple and the striking of the match?

Male: Yes.

Lecturer: (Slide 38) You would? Okay, so some people think that just the correlation is sufficient to give us a causal relation, other

people would say, ‘We want to see whether it makes sense, first.’

We can’t think of a mechanism, can we, whereby the dropping of a pineapple from a tree would lead to the striking of a match? Of course, that may be just our lack of imagination.

Male: It goes back to all the swans.

Male: An example of that is that effectively, very few people have a sneaky view sunspots caused stock markets to go up on [inaudible] for wherever. Effectively it worked out that sunspots then therefore caused subsequent good harvests, and therefore grains came up and there was prosperity. So, it's a linking point, and people just didn’t understand what those linking points were.

Lecturer: Well, and the same could be true of pineapples and striking matches?

Male: Yes.

Lecturer: We need to have a story told, don’t we, before we start to find it convincing?

Thinking about causation, nearly everyone knows someone who smoked for decades. So my father, he died of emphysema so my father is not a good example, but most people know someone who has smoked for decades, and yet

has no sign of lung cancer in their eighties. Does this mean that smoking doesn't cause cancer, do you think?

All: No.

Lecturer: Does it?

Female: No.

Lecturer: In exactly the same way, you say the sighting of one swan doesn't really show that the argument, 'all swans are white' wasn't a good one. It shows that that conclusion is wrong.

In the same way, actually, here we know that we've got the natural uniformity wrong somewhere, haven't we? So, if you make a distinction between the *natural necessities* that we think of as ruling our worlds, the laws of nature that exist and are as they are, quite independently of us. That we try and discover, we try and observe, so there are *laws* which are linguistic entities.

We're trying to get our *laws* to correctly describe these *natural necessities*, aren't we? We think whereas lack of correlation is evidence for error, isn't it?

If we think A causes B, and then we see an A and not a B, that's evidence for error there, isn't it?

Male: Yes.

Lecturer: What we know now is that it's not that A causes B, it's that *A** cause B, or that As in certain context cause Bs. Or that there's something, in addition to A, that causes B.

So, I think that what we know, from the fact that we do know people who haven't developed lung cancer, is that there is something more to the claim that smoking causes cancer, than that just smoking causes cancer. Because if it was just smoking, there wouldn't be any smokers who didn't have cancer.

Female: Nor does it say inevitably, does it? It doesn't say, 'Does this mean that smoking inevitably causes cancer?' which would clearly be wrong, judging by the first sentence, but it doesn't mean it doesn't cause it.

Lecturer: Yes. We think of causation; if A causes B, we think of A as being *sufficient* for B. Therefore, if you get an A without a B, you know that A isn't sufficient for B, and you need to add something else to it.

So, we know we haven't yet quite got there. There's something more to it, maybe there's a genetic element, or an environmental element, or something like that. We do know that it can't be smoking just on its own, but we do know that actually the probability of getting lung cancer from smoking is very high. Which is, as you say, it's not deterministic.

One or two (questions), and then we'll move on.

Female: A causes B; does it mean that A is the difference? Is this to be interpreted that A always causes B?

Lecturer: If we say A causes B and then we see an A and not a B, that's (the sighting) usually thought of as falsifying that (the hypothesis), isn't it?

Female: Is it?

Lecturer: I think that most people, I don't know if this is a language thing, but I would have thought that most…

When a scientist is trying to show that A causes B, what you're trying to show is that A is *sufficient* for B.

Female: You could take it for your example, because a birch pollen causes asthma. We certainly don't believe that everyone, it depends, you have to have an allergy first.

Lecturer: Well then you know that birch pollen doesn't cause asthma, there's something in addition.

Female: [inaudible] the person, yes, but in people who are susceptible to it.

Lecturer: Yes, there's something in addition, as well. You could take the word 'causation' to mean a probabilistic type of causation, and I agree that in many cases we do mean it thus, I think.

So when we say that smoking causes cancer, we don't mean it's deterministic, but we do mean there's a high element of risk of getting lung cancer, if you smoke. Very high element of risk.

Male: It could be that smoking always causes lung cancer, but sometimes you don't live long enough for it to happen. So, the limiting parameter could be age. You might die of something else before it actually…

Lecturer: Yes, so if this chap went on to live to 105, he might get it too?

Male: Yes.

Lecturer: Yes, could be.

Male: Are there any examples of lung cancer without smoking?

Female: Yes.

Male: Is it high incidence? I wouldn't know.

Lecturer: I wouldn't know, either. Does anyone know?

Female: I do.

Lecturer: Go on, is there high incidents of lung cancer without smoking?

Female: There are causal things, if you work with asbestos and you smoke you will get lung cancer. Or, [inaudible] it's as you said, it isn't exclusive, or determinate, it's putting it together.

Lecturer: Also, passive smoking; having grown up in a house where somebody smokes 60 a day for 16 years - mind you, I smoked myself I have to admit, at one point, as well, but let's say I hadn't - then I might still get cancer that was still smoking linked, even though I had never myself smoked.

Anyway, this is a red herring, let's carry on. You can do these at home. 'Sleeping with your (Slide 39) on gives you headaches.' Is that a good causal argument?

If you've seen a correlation between sleeping with your shoes on and headaches?

Female: Yes, I can see that, people being drunk.

Lecturer: So there might be a common cause that leads to both sleeping with one's shoes on, and having a headache? Okay, do that at home.

(Slide 40) Evaluating analogies; 'you greatly enjoyed Barchester Towers, you I'm sure you enjoyed Can You Forgive Her? They're similar in all sorts of ways'. Okay, is the claimed similarity relevant to the conclusion drawn?

If you say, ‘well, in what way are they similar’, and I said, ‘well, they were both published by Penguin’, or, ‘they’re both in green covers’, you might think, ‘what has that got to do with it?’

We might want to know how many respects they’re similar. So, if I tell you it would be a good thing to buy a Toyota Yaris, because mine has been fine, then the fact that it was also owned by a very safe middle aged woman for many years, and she looked after it very carefully and all these things. The more I can add similarities between the one you’re thinking of buying, and the one I’ve got that I’m prepare to recommend, the better the analogy.

We might want respects in which the two things are dissimilar. So, the subject matter in Barchester Towers differs from Can You Forgive Her? If you really only go for ecclesiastical novels, then the fact that this isn’t one, or at least not as obviously, might make you think, ‘well, actually that’s a disanalogy that would matter.’

As with the watch-maker one. Are there disanalogies between the universe and the watch that make us think that even if a watch does have a maker, maybe the universe doesn’t?

One question about that.

Female: I just wanted to say that there are far less. I don’t see any relevance to the watch-maker and the universe thing, whereas there is a relevance, here.

Lecturer: With the watch-maker and the universe, the idea is that the watch is regular and runs according to a mechanism. In the same way, the universe is regular, runs according to mechanisms.

Female: I see.

Lecturer: So that's the idea.

(Slide 41) In 1992, the marketing people at Hoover reasoned that because the take up of air tickets to Europe on a recent promotion had been low, a promotion offering air tickets to the US in exchange for buying £100 worth of Hoover products would also be low. Was it?

No, it wasn't and actually they lost a huge amount of money as a result of this bad analogy. What are the disanalogies between air tickets to Europe and air tickets to the US, on the basis of buying £100 worth of products?

Female: Because tickets to the US are far more expensive.

Lecturer: Exactly. Actually, this was a hugely cheap way of getting to the US, if you only had to buy £100 worth of products, you were, in effect, getting to the US for £100. They lost certainly millions, I'm not sure if it wasn't billions.

(Slide 42) So, there you are, some analogies to evaluate at home.

(Slide 43) Evaluating arguments from authority; several times in these lectures, I've told you you should just believe me about something. Well, should you?

Male: Of course.

Lecturer: Thank you. Is the authority qualified? Do I really know what I'm talking about, here? Is the authority qualified in the appropriate area?

Actually, there are many more philosophers far more qualified than I am in logic and critical reasoning. So, if you have another authority who disagrees with me, then you should probably believe that authority, rather than me.

Do other authorities agree? Are there other philosophers who disagree with Marianne? Well, I can tell you right now there are many, many, many of them, but probably not on the things that I told you you should take from me. Because on the whole, I would have explained that there was disagreement, if that was the case.

Female: Apart from earrings and jeans?

Lecturer: Yes.

Okay, could the authority be biased? Am I being paid by somebody who wants me to carry on reasoning appallingly, to mislead you? It depends how much they offered me.

Can references be provided by whoever appeals to the authority? So, for example, can I provide references for the claim? A couple of times, I hope I've demonstrated I can.

Also, if somebody said, 'Well Marianne said that, therefore it must be true to you', you would be within your rights to say, 'well which lecture was it she did that? Can you show me on the handouts, and so on?' We need references when people appeal to authorities.

Male: Should be telling this to the BBC.

Lecturer: (Slide 44) ‘Judy Dench, Sting and Julie Christie urged David Cameron to decriminalise possession of drugs.’ This was in all the papers, presumably because of Dame Judy Dench and so on. Is this an appeal to authority, do you think?

Male: No.

Male: No.

Male: Celebrity, it's rubbish.

Lecturer: It's a sort of appeal to authority, isn’t it? Otherwise, why introduce Judy Dench?

Male: It's an appeal to influence, isn’t it?

Lecturer: It's an appeal to influence, that’s right. Does it work? Should we listen to Judy Dench on drug policy? No? Even though we love her and we think she’s wonderful and we really enjoy her films?

Male: It depends on your sample.

Female: I think it depends on basing the thing on knowledge.

Lecturer: We have no reason to think that Judy Dench knows any more about drugs than…

Female: Some people, just not Judy Dench isn't it? They are convinced by Judy Dench, it depends the level of the [inaudible] of the people that do.

Lecturer: We're running late, (Slide 45) those two answers are just to do…

(Slide 46) Finally, walking along the beach, Fred, Hilary and Sam - and those of you who know anything about Hilary Putnam will know that that's who I am referring to, here - came across an excellent likeness of Winston Churchill drawn in the sand. Hilary expressed amazement that crawling ants could have made such a picture. Sam said, 'There's no way it could be ants, it must have been the wind'. Fred insisted it must have been drawn by a person.

We might say, well obviously we should assume Fred's explanation is the true explanation, it's the best of the three. By 'best', we mean it's the most conservative one. It's the one that's most consistent with our other theories, Etc.

Here, we need to ask, 'Do we have other hypotheses?' We've got three, here. Could we offer another for why this likeness was drawn in the sand? Might there be an even better hypothesis? If saying that a hypothesis is the best, but it's the best of a bad lot is not much use, is it?

Are there hypotheses we're unable to formulate? So, for example, some people believe that it's obvious that science is going to be able to explain all sorts of things like God, morality, rationality, etc, etc. I sometimes say to them that this is like a monkey being half way up a tree on the way to the moon, and it says, 'so far so good!' You think, 'actually, there's a principled problem, here. How do we know that we're not in this situation with respect to this type of explanation?'

(Slide 47) That was just an example of a use of abduction in science, because science uses abduction a lot, as you can imagine.

(Slide 48) There's an abduction to evaluate at home, and (Slide 49) there are some references for all the things that I've talked about today.

(Slide 50) Next week, we're going to be doing fallacies. It's the last week, all of you have brought stuff for me to have a look at, and I'll try and have a look at it before next week.

I'll also devote some of next week to answering your questions. So, if you have questions, can you bring them along to next week? I'll try and answer them then.

Sorry we've overrun.

END AUDIO