

## A romp through the foothills of logic – Session 5

Lecturer: What would people like to talk about?

[inaudible discussion]

Female: (P arrow Q)?

Lecturer: The meaning of (P arrow Q)? Let's get the truth table up. It would be nicer to have the ... ignore everything except this. That's the truth table for the conditional.

What Kirsten wants me to do is to motivate that – the meaning of that – because we don't usually think of a conditional as being true in cases where – we can see that it's reasonable to say that 'If P then Q' is false, where 'P' is true, 'Q' is false. I think that's a reasonable thing, isn't it?

But to say that 'If P then Q' is true, just because say here – where 'P' is false and 'Q' is true – you might think, 'Well, hang on. If the postman is...' I have to think of examples here. I hate doing that. 'If it's raining, you will get wet.' Well, why should that be true if it's not raining? Well, one answer to that was given to me by Steve a moment ago. Do you want to try again? You were right last time, so let's tell everyone else.

Steve: My assumption was that in that example you can get wet in some other way, so if it's raining is false, if wet is true. I think

there's a hidden axiom there, that unless something is obviously false then you can count it as true.

Lecturer: Well, bivalence tells us that unless it's true – the 'obviously' doesn't come into it – if it's not true, it's false.

Steve: Yes.

Female: But somebody could have thrown a bucket of water over you.

Steve: So therefore you count it as true.

Lecturer: The thing is that it's true, isn't it, of my finger, that if it were a duck – well, perhaps I will say a sparrow. If it were a sparrow, it could fly. My finger isn't a sparrow, but if it were a sparrow, isn't it true that it could fly?

Male: Yes.

Lecturer: So that motivates that line of the ...

Think of it like this. 'If P then Q' is the same as 'it's not the case P and not Q', is that true? So 'if P then Q', in English, means the same as 'it's not the case P and not Q'. Fair enough?

The truth conditions, if I draw the truth table for that, I get this: 'not-(P and not-Q)'.

That's the ampersand. Again, it's just my computer prefers an ampersand.

So, 'not-Q' just reverses that ('Q'), so we get false, true, false, true. 'P' and that gives us – so let's put the small ones in here under 'P'. So that and that is only true if they're both true. So here it's true and otherwise it's false.

Are you with me? The not reverses that line, so we get true, false, true, true, which is the same as that. Do you see? True, false, true, true, true, false, true, true.

So the truth conditions of 'if P then Q' are exactly the same as the truth conditions for 'it's not the case P and not Q'.

Female: Is it in the nature of a causal ...?

Lecturer: Absolutely. That's the trick of realising this. All the sentence connectives, all the truth-functors, relate completely independent sentences. There is *no relation* between the two sentences related. But when we use 'if P then Q', in English, it's nearly always the case that we use it thus, because we think that P and Q are related in some way -

Female: One causes the other?

Lecturer: Well, not necessarily causal. I might say, 'If there's a double yellow line you shouldn't park on it.' Well, that's nothing to do with causation. It's to do with the legal situation.

There are lots of times where you have a relation that's nothing to do ... 'If it's a sibling, and a sister, then it's female.' Again, that's not causal. That's a different sort of relation.

So it doesn't matter what the relation is. When we use 'if P then Q', in English, we tend to use it only when there's a relation between the two, but logic sees no relation between the two. You're importing your understanding of Gricean conversational rules, believe it or not. (Laughter)

Male: I think the girl was saying earlier that in your example, 'If it's raining you will get wet', if it's not raining, it doesn't falsify the fact that you're still wet.

Lecturer: That's right, because there are other ways. Somebody might have thrown a bucket of water over them.

Male: So it's not false, therefore it counts as true?

Lecturer: Yes.

Male: So it counts?

Lecturer: Yes, that's right. The fact that it's not raining doesn't mean that I won't get wet, because somebody might have thrown a bucket of water over me.

But that doesn't change the fact that if I'm saying, 'If it's raining I will get wet' means that 'It's raining' is *sufficient* for my getting

wet, but not *necessary*. It would only be necessary if it was a *biconditional*. It would only be necessary *as well* if it was a biconditional.

Male: Earlier on you talked about the \_\_\_ connective 'because'. No, it wasn't the \_\_\_ [cross talking]. Sorry, the connective, anyway, 'because', and showed that there was not a truth functional. 'If P then Q' looks like 'Q because P', but it's not.

Lecturer: It's not. It's very definitely not.

Male: That's I think what confuses.

Lecturer: Well, again, you're importing the Gricean Conversational Rules into the ... I could do this quickly for you. I think I could do it quickly. Shall I try? (I'm not sure actually. Let's move that over so I've got more room).

Female: Why do they even have the arrow? Why not just have and, and not P?

Male: You could make them all up from just 'and' and 'not'.

Lecturer: Yes.

Male: Any of them and not.

Lecturer: All of these truth-functors can become a combination of 'and' or 'not'. They can all become a combination of 'if' or 'not' as well. So we only use the five truth-functors, because actually it makes things quicker, but you could actually carry down to two if you wanted.

The Gricean Conversational Rules – sorry, Chris, you will have heard of all this.

Chris: I can't remember it though!

Lecturer: Oh, good.

Chris: A second attempt would be good.

Lecturer: Phil asks me whether Fred is a good philosopher, and I say, 'his handwriting is excellent.' What do I *mean* when I say that?

[inaudible discussion]

Lecturer: Yes. So what I'm saying is, 'he's a bad philosopher', aren't I?

Now, that's interesting. Here we have a sentence with a particular meaning. That's not the meaning of the sentence, *that he's a bad philosopher*, but I've used that sentence to say, 'he's a bad philosopher.'

Well, how did I do that? The reason I was able to do it is because we've got to distinguish *sentence-meaning* and *speaker-meaning*.

Sentence-meaning is the strict and literal truth conditions of the sentence. So the strict and literal truth conditions of the sentence, 'His handwriting is excellent', if you break that down you're applying a predicate, 'is excellent', to the subject, 'his handwriting'.

So that sentence, if I write it here, 'His handwriting is excellent', I'm not *using* that sentence. I'm just *mentioning* the sentence, if you like. I'm putting quotes around it. I'm mentioning it.

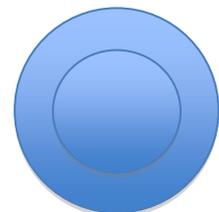
I can say, "His handwriting is excellent' has four words.' Because I'm talking about the sentence, not about his handwriting, once I put quotes around it.

I can say, 'The meaning of that sentence is that we're saying his handwriting is in the class of things that are excellent.' But I didn't mean that, did I, by the way I used that sentence?

So there's a different thing called *speaker-meaning*. The weak meaning of that sentence is 'His handwriting is excellent.' The strong meaning, i.e. the meaning in context, was, 'He is a lousy philosopher.'

[I used this diagram to explain meaning:

The innermost circle represents the strict and literal meaning of a sentence, the next circle the force (interrogative force, imperative force, or assertoric force...) with which the sentence is uttered. Then you've got the mood (sarcasm, anger...), the next the context in which the sentence is uttered.]



Here's the strict and literal meaning of something, at the very centre of something. On that you've got the force of the sentence, so whether you use it with interrogative force, or imperative force, or assertoric force. Then you've got the mood. So if I say [in a very angry tone] 'I'm not angry!' Sorry. (Laughter) What have I said?

Female: You're very angry.

Lecturer: Or rather I have *implied* that I am angry rather than not angry. So it's completely the reverse.

What's your name?

Male: Tony.

Lecturer: Tony? So, this is the tenth of a class of ten classes, and every single time Tony has come in ten minutes late. The door slams open and Tony rushes in and sits down. The tenth one he comes in like that and I say, 'Hello, Tony, early again?' See, I've completely reversed it.

So the mood of sarcasm, or anger or whatever, can operate on the sentence to reverse its meaning, and the context of course is a very important part of this.

You wouldn't get that if this was the first time you've been to the class; except you probably would, actually. You would probably realise that I was implying that Tony was always late, rather than saying that he was early again.

There is somebody called Paul Grice, who gave us three rules of conversation or maxims of conversation. The idea is that if you're conversing with somebody, and that's what you're trying to do, you're trying to obey all the rules in such a way that you can communicate as easily as possible, that these are the rules that you've got to obey.

The first one is 'Be truthful.' So I shouldn't say 'P' unless I believe P. So it's violating the rules of conversation to say 'P', when actually I don't believe P. So if I assert 'P' when I believe P to be false, I'm not playing the conversational game, am I? I'm trying to con you in some way.

Female: Unless you're in a context where it is requested of you to say that, even though you don't believe that.

Lecturer: Well, in that case I'm not asserting P, am I? If I'm in a play and I say, 'To be or not to be', which isn't a ... If I utter a sentence in assertoric mode, nobody in the audience actually thinks I'm expressing a belief, do they?

Female: They would in a debate. In a debate you might be expressing the opposite view of what you believe.

Lecturer: That's true, but then I would *postulate* P. Or I would say, 'Let's, for the sake of argument, believe P.' But if I actually assert P, believing not-P, I'm not playing the right game. I'm lying to you, aren't I?

Male: Or being courteous. 'I'm pleased to meet you.' (Laughter)

Lecturer: Well, it's still a lie, isn't it?

Male: 'Thank you for a lovely time.'

Lecturer: It's a white lie.

Male: Oh, it's a lie, but it's a courtesy.

Lecturer: Okay, but it's still a lie. We're not talking about courtesies at the moment. But it's true that I'm not playing the conversational game.

In other words, what I'm not doing is communicating meaning here. I'm just going through the motions of a convention, aren't I? 'Pleased to meet you' or whatever. So being truthful is the first one.

The second one is 'be informative', or be as informative as possible, and no more than necessary.

This diamond means 'possible' and that square means 'necessary'.

If you say, 'Where's my wife? Have you seen my wife?' and I just saw her in the kitchen kissing Fred, I don't want to say, 'she's in the kitchen', because he's going to hare off there and find her. What I'm going to say is, 'she's in the kitchen or she's in the bathroom.'

Now, I haven't actually lied to you, have I? But nor have I played the conversational game. I haven't given you as much information as I *could* have given you, because I could have told you, 'she's in the kitchen', which is where she is. I've given you less information than actually I could have done.

In the same way, if I give you more information – you've heard of kicking sand in someone's face – you ask me a question, and I don't want to answer it, so I give you a blizzard of answers. So there I'm giving you more information than I need, but again I'm not playing the conversational game.

What's the third one? Who can tell me what the third one is?  
What's the third one?

Male: [inaudible]

Lecturer: Oh, I know, 'be relevant'. If you ask me whether he's a good philosopher, and I say, 'His handwriting is excellent', my violation of the third Gricean maxim of conversation immediately alerts you to the fact that I don't mean what I'm saying. I'm saying something *other* than the sentence I'm using is saying.

Female: You started off saying it was two things about the speaker, and implicit in your scenario is a hearer. The hearer has to be there, together with certain rules which they understand, so that they can gain ...

Lecturer: The Gricean [account of meaning] is the speaker-meaning 'I mean something when I intend you to understand something

by what I say and I intend, in uttering whatever I say, to want you to understand what I believe, and as a result of your coming to believe that I believe whatever it is I'm saying you, too will come to believe it.'

So when I say, 'his handwriting is excellent', I intend you to believe that I believe he's not a good philosopher. What's more, I intend to make *you* believe he's not a good philosopher because *I* believe he's not a good philosopher. So those are the Gricean rules of meaning.

Male: We got into horrible trouble last week because some of our class didn't have English as their first language. So you ended up, if you've got that flipchart, on that outer ring you also put cultural in, if I remember correctly.

Lecturer: Well, I think that cultural comes under context. A French woman who was in the class said that my saying, 'His handwriting is excellent' would not have meant to her, 'He's a lousy philosopher', which I find really interesting, because it's so obvious to us, but that's ...

Male: There was a Danish speaker as well or something that had a similar view.

Lecturer: There was only one that I remember.

Male: [Yes, I think].

Lecturer:

How this links with that is this. Let's move this away now.

If we're going to stick to the first rule, you wouldn't say 'If P then Q' if you knew that we were in the world in which 'P and not Q'. So there's absolutely no reason to utter 'If P then Q' if you think you're in this world.

If you think you're in *this* world, where the antecedent is false, why would you say, 'You're going to get wet'? Sorry, why would you say, 'If it rains you're going to get wet', when actually you know that you're in this world, and you're going to get wet whether it's raining or not, because I've put a bucket of water above the door? (Laughter)

There's no reason to utter 'If P then Q' if you believe yourself to be in the world in which not-P.

It's the same down here. If you believe yourself to be in the world where it's not raining, why would you say, 'If it's raining you're going to get wet'? You would just say, 'You're going to get wet', or, 'It's not raining.' So you would be obeying the second rule of conversation.

So there's no reason to utter 'If P then Q' here (the first line of the truth table. There's no reason to utter 'If P then Q' The third line of the truth table here. There's no reason to utter 'P then Q' here (the fourth line of the truth table).

If you know that it is raining *and* you're going to get wet, why don't you just say, 'You're going to get wet'? You don't need to say, 'If it's raining you're going to get wet.'

So there's absolutely no reason to utter 'If P then Q' [because you wouldn't utter it in the world that is the second line of the truth table because in that world it is false. MT], so why do we use 'if P then Q'? Answer: whenever you have reason to think

there's a *relation* between P and Q, such that *it is not the case P and not-Q*, so you say 'If P then Q.'

It might be a *legal* relationship. It may be a *conventional* relation. So, 'If you're going for coffee go to the Acland Room', and there's a relation between those two things that is just true in Rewley House today, isn't it? The carpet will be dry by Monday [coffee is usually served in the common room]!

Female: If you did want to let the relationship exist as an entity, would the last two then become a third type of value which is invalid or ...

Lecturer: No. When you say the last two, you mean these two?

Female: Yes, where you have 'not' ...

Lecturer: No. The fact is these are the strict and literal truth conditions of 'if P then Q'. What I'm giving you now is the speaker-meaning. The *sentence-meaning* of 'if P then Q' is given by that [the truth table].

When we use it in conversation, it's the speaker-meaning that's important, and the speaker-meaning of 'if P then Q' will nearly always imply a relation between 'P' and 'Q', whereas the logic, the strict and literal truth conditions, there is no relation between...

Which is why I can say 'If P then Q' even when the relation is one that lasts for two seconds, or the relation is only one that I put in there.

It doesn't have to be a real relation in the world. It doesn't have to be a causal relation. It doesn't have to be any sort of particular relation.

It's just that we wouldn't use 'if P then Q' unless we have reason to think there's a relation, but that's *not* part of the sentence-meaning. It's only part of the strict meaning.

Female: So, in speaker meaning you'd just drop off the last two?

Lecturer: Well, no, the speaker-meaning is not given by the truth table at all. The truth table only gives the sentence-meaning.

The truth tables are all in here, and we take something with that definition and we apply a force to it, we apply a mood to it, we do it in context. We use the sentence in a certain way, and that gives us the speaker-meaning. That's the weak, strict, and literal meaning of the words 'if, then'.

Male: But the [inaudible] is an odd one out, because it's the only one that seems to strike me as being based on probability.

Lecturer: Well, it's *not* based on probability. That's what I'm just explaining.

Male: Well, say for example, 'If it is raining, then you will get wet.' No. You *might* get wet, if you didn't take certain ... There's a probability here, where if you don't do something...

Lecturer: But that's irrelevant. The sentence-meaning, the strict and literal meaning of 'if P then Q' has no relationship in it at all – probability, causation, any other sort of relation.

Male: If I accept that, then I have a problem with the third one down ... I have a problem with that, because I can't get ...

Lecturer: If my finger were a sparrow, it could fly. Is that not true?

Male: No.

Lecturer: Isn't it?

Male: No, because it might be a sparrow with a broken wing.  
(Laughter) It's not *necessarily* true.

Lecturer: Okay, if my finger were a *normal* sparrow, it could fly.

Male: I'm glad you used the word 'could' because that implies a probability.

Lecturer: You're missing the point of what I'm saying here.

Male: Well, I may be.

Lecturer: Yes, completely.

Male: I don't have a problem with accepting the rule, and always applying it, but I'm ...

Lecturer: Which rule are you accepting?

Male: If 'P' is false and 'Q' is [inaudible] then 'P arrow Q' is always true isn't it? It's true. That's a rule, and I can accept a rule.

When we start talking about some of the layers beneath it, whether it's the speaker-meaning or the sentence-meaning, whatever, I'm looking to consequents. You're saying you should be able to divorce it from consequents.

Lecturer: What I'm saying is that the strict and literal meaning of any sentence constituted of two sentences – whatever the sentences are – and the conditional 'if then', the English conditional 'if then'.

So actually forget that, just say 'if P then Q' is given by the truth table definition, the strict and literal meaning, and that isn't really a meaning at all.

So if I write 'Lucy is fat' here, you all understand that, but you couldn't determine a truth condition of that, could you, because you don't know who Lucy is?

I'm not actually talking about any Lucy. That's actually a *type* of sentence that I could use to express a belief.

If I used it with its strict and literal truth conditions, to express a belief, the belief I would be expressing is that somebody, Lucy, is in a class of things that are fat.

So that has the weak meaning only. It only has strict and literal truth conditions. It doesn't have anything more than that.

But I could use that sentence to say *anything* I choose. I could use it to say, 'Okay, we're going to have a convention in this class. At 12:30, when this class ends, I'm going to say 'Lucy is fat', and what that means is that the bar's open and I want a glass of red wine.' (Laughter)

So now 'Lucy is fat', the very same sentence, which continues to have the same sentence-meaning it always had, the speaker-meaning of that sentence in the context of this class is completely different, isn't it, and much more interesting?

So we're distinguishing speaker-meaning and sentence-meaning. I'm saying that the strict and literal meaning of 'if then' is given by that, but whenever we use it in language it's to say that there's a relation between P and Q.

Now, that relation is not given in the truth table. That relation could be one of probability. It could be one of causation. It could be one of law. It could be one of convention. It could be any sort of relation at all. But that relation is not part of the strict and literal meaning. That's given by that [the truth table].

Female: Can I ask a different question?

Lecturer: Nothing to do with this? Is yours to do with this?

Male: No.

Female: Can I ask a different question sometime?

Lecturer: Yes. Well, we can move on from this. Do you want to ask a question?

Male: Yes, I do have a question to do with this. Maybe if we go back to you saying 'Paul Grice is not a good philosopher.' You were saying that was your intention. However, you say 'His handwriting is marvellous.' It may be a professional clash between two philosophers.

Now then, you could be right. Paul Grice could be right. You could both be good philosophers. However, he's got on your nerves because he's said something you think is rubbish, but you still think he's a great philosopher, and he still thinks you're a great philosopher.

Lecturer: This is irrelevant. What I was doing was explaining the meaning of Fred's... – not Paul. Paul Grice is a brilliant philosopher. (Laughter) It's *Fred* whose handwriting is lousy.

I'm implying, even though I'm not saying, that Fred is a lousy philosopher. I may be wrong. My belief may be false. But I'm explaining the *meaning* of Fred's handwriting, not the truth value of it.

So, different questions. You had a different question?

Female: Oh, yes. Could you please explain two things you referred to yesterday? You said you would explain today the principle of the uniformity of nature and paradoxes of confirmation. You said, 'I will explain them tomorrow.'

Lecturer: Oh, jolly good. These are nice questions. The principle of the uniformity of nature and the paradoxes ...

Female: That spells 'p-u-n'. I love it.

Female: It's great.

Female: That's paradoxical.

Lecturer: Do you know, I've never noticed before, but that spells 'pun'.

Female: Yes. That's what I was saying.

Lecturer: I've never noticed that before.

Female: I think that's a wonderful paradox.

Lecturer: I look at it and I read, 'Principle of the Uniformity of Nature.'  
(Laughter) Isn't that interesting?

Female: Absolutely.

Lecturer: So, Principle of the Uniformity of Nature.

The philosopher David Hume noticed that we reason in two different ways. He talks about there are arguments from fact and there are arguments from reason.

Arguments from *fact* always take the same sort of pattern. So we say, 'Every time I've seen a daffodil it's been yellow. Therefore, the next time I see a daffodil it will be yellow.'

Actually, this is a bad example, but never mind, bear with me.

Female: Inductive reasoning, isn't it?

Lecturer: Reasoning from that, as a fact, is inductive reasoning, exactly.

Notice that what I do is I go from having observed a correlation in the past to assuming that I will observe the same correlation in the future. If I see A and B, and A and B, and A and B, I will *induce*, I will reason, that next time I see an A I will see a B, and so on.

If you think about it that underlies everything we do in science. If I manage to produce something in my lab, and I think, 'Oh, that's interesting, so 'if A then B'.' Notice we've immediately got a causal relation, or a probability relation or something in here, 'if A then B'.

I ring you up in Australia, and I say, 'Can you just check that you too get this? Is it the case that if A then B for you? Is it the case that if A then B for you?'

It's only if you get an A without a B then it stops looking like a causal relation, doesn't it?

Say I see lots of correlations, and I form a hypothesis 'If A, then B.' I need to try and falsify that hypothesis, and I can do that if I find an A and not a B. So I'm going to try and bring about an A without a B.

Now, if I succeed, I haven't shown – well, I've shown that it's not the case if A then B. There's something not quite right about that.

Now, I may be so sure of my hypothesis that what I start thinking is, 'Okay, it's not the case that it's *always* a B, but I bet A\*s are Bs. There is something about A. It may be that there are some As that don't have a B, but most...'

For example, 'smoking causes cancer.' Well, we all know that there are smokers who don't get cancer. So what we know is that the law 'smoking causes cancer' is not right, but we're still pretty sure that there is something about those people who smoke, that if we could find it we would know that people with gene A, who smoke, will get cancer, or whatever.

We feel so sure about that generalisation that we feel that the law could be tightened in some way, if we can find a way of tightening it. So to falsify something is not necessarily to think of it as false, but it's certainly to think of it as in need of being tightened.

Underlying all of this is the assumption that the future is going to be like the past, so if smoking has caused cancer in the past, smoking will cause cancer in the future.

Actually, if you think of it, *all* our inductive reasoning it's based on the Principle of the Uniformity of Nature.

Why do you think the Sun will rise tomorrow? Because it's risen every day in the history of the universe. Why do you think your cat is going to scratch you when you go home tonight? Because it scratches you every night when you go home. And so on.

We cannot reason without resting on the Principle of the Uniformity of Nature, but nothing we do can justify that belief. We all believe that nature is uniform, but we can't justify it. We can't justify it *deductively*, because no contradiction is generated by assuming that nature isn't uniform.

Do you see what I mean? You can't get a 'P' and 'not-P' from assuming that nature isn't uniform. You can't justify it *inductively*, because then you're arguing in a circle.

Why do you think the future is going to be like the past? Well, because in the past the future has always been like the past, in the past. Therefore in the future the future is going to be like the past, in the future.

Do you see? You're just arguing in a circle. All inductive reasoning rests on the principle of the uniformity of nature, which cannot be justified.

Female: Isn't it just probability? You're saying it's *highly* correlated, it's *highly* probable, but it's not perfectly caused.

Lecturer: Well, as I say, you can never get certainty from inductive reasoning. The reason you can't is because all inductive reasoning rests on something that cannot be justified.

Female: But do you need it to be justified? Rather than cause up there, can't you just say correlated?

Lecturer: It's very rarely the case. We don't think that causation necessarily involves determinism. Sometimes it does. For example, 'smoking causes cancer.' We tend to think that that's true, but we know that, to the extent that it's true, it's not deterministically true.

Female: Not in the individual case?

Lecturer: Yes, that's right. 'smoking probably causes cancer' is a much more accurate statement than 'smoking causes cancer', but that's because we know that there are cases of smokers who don't get cancer.

Female: If you add that 'probability' into it, we can get certainty on the view of probability.

Lecturer: Well, that's a very interesting one. We can be certain that smoking probably causes cancer, but does that actually add anything?

Male: I'm not sure I would phrase it that way. I think I would phrase it as 'If 100 people smoke, 14 of them will get cancer.'

Lecturer: You can be much more certain with something like that, because you can check it. Epidemiologically you can say, 'that's what we've found.' But do you see you're still resting on the Principle of the Uniformity of Nature? In the past studies have shown that for every 1,000 people, did you say, 14 of them...

Male: Every hundred. It's about one in seven.

Lecturer: For every 100 smokers, 14 of them will get cancer. Is that all?

Respondent: Well, if you compare that to not smoking, where it's about a half in a hundred. One in seven is thirty times greater than not doing it.

Lecturer: But even so, it's less than I thought ... (laughter)

Male: If I handed you a revolver, with one chamber left, you would be loathe to pay me for the privilege of pointing it at your head and pulling the trigger.

Lecturer: You're absolutely right. (Laughter)

Male: Lovely.

Female: ... If you actually said, 'Based on the statistical analysis to date we have seen X cases, and here is the number of...'

Lecturer: 'You will certainly probably get cancer.' (Laughter) How does that help you?

Male: But it's the number to date that you mentioned. That's the key thing. It's the number *to date*.

Female: If you actually bypass relying on the principle and state it in your own conjecture, and say, 'To date we have seen X cases. This number has been ...'

Lecturer: You could say that.

Female: You then have a relationship ...

Lecturer: You're absolutely certain that in the past, out of every hundred people who have smoked, 14 of them had cancer.

Female: Had cancer.

Lecturer: Had cancer, yes.

Female: But if you just state the fact ...

Lecturer: Of that you're certain. If you then move to the future and say, 'Therefore, of every 100 people who smoke in the future, 14 of them will get cancer', you're extrapolating to something of which you're not certain.

The reason you're not certain of it is because you're relying on the Principle of the Uniformity of Nature. You're assuming that the future will be like the past.

Female: [inaudible] assume it, if you don't make the [cross talking].

Lecturer: You've got to assume it. You can't not assume it.

Female: No, I mean if you just have a way of stating correlation [between] things. Basically you just put the stats out there, and you're leaving it to the other person to assume.

Lecturer: So you're putting the stats out that you've observed in the past, so that people can look at them and make their own decisions?

Female: Yes.

Lecturer: So what are they deciding about?

Female: It could be anything. So they could see the information of 'up to date maybe we didn't have enough methods, or in the future they could have a piece of information, that can be used, that doesn't make assumptions.'

Lecturer: But one of the reasons you want it put out there is so that people can decide whether or not they should smoke.

So they might look at it and think, 'Ah. Wow! That's like putting a loaded gun to my head, even though most chambers are empty. I'm never going to touch another cigarette.' Or something like that.

If you make any sort of decision about your actions, on the basis of that correlation you see in the past, you are relying on the Principle of the Uniformity of Nature.

If you as an epidemiologist, or as a researcher into cancer, design your next experiment on the basis of those facts, you are again relying on the Principle of the Uniformity of Nature.

You cannot in inductive reasoning get away without relying on the Principle of the Uniformity of Nature. So anyone who tells you that scientists don't rely on faith, they are completely wrong. They don't rely on faith in God, but they rely on faith in the Principle of the Uniformity of Nature.

Male: Could I just ...

Lecturer: If they're hypothesising.

Female: No, I mean if they're not doing either of those things, and they just want to share information ...

Lecturer: Yes. If all they're doing is saying, 'This is the case.' It's when they say, 'This will be the case', or, 'I'm going to assume that this is the case in order to...' But how many scientists don't do any predicting?

Female: Hopefully some good ones.

Male: Can I just ask? You can say, 'Yesterday the Earth revolved around the Sun in a fixed orbit. Today the Earth is revolving around the Sun in a fixed orbit, and the Earth is spinning on its own axis at the same time.

Therefore, it is reasonable to say, 'Tomorrow the Earth will continue to revolve in its own axis, will continue to go in a fixed orbit, that we know about, around the Sun. Therefore, from our perception the Sun will rise, all other things being equal.'

Lecturer: And that's only reasonable if it's *reasonable* to rely on the Principle of the Uniformity of Nature.

Male: No, I'm standing ...

Lecturer: I've just said that this can't be justified.

Male: There's no hidden heavenly body approaching to the Sun.

Lecturer: Let's move on.

Male: The thing that made an impact for me was Russell's example about the chicken who comes down each morning. The chicken comes out of his den and the farmer has fed him. So he comes out on Monday morning and expects the farmer to feed him, and the farmer wrings his neck.

Lecturer: How do we know we're not in the position of Russell's chicken? That's right. Turkeys and Christmas is a perfect ...  
(Laughter)

Female: Could you answer the second part?

Lecturer: Yes, in a minute. We just have a bit to go on this.

Male: No, I was going to ask you to derive  $(P \vee Q)$  using only 'not' and 'and'.

Lecturer: 'P or Q' using only 'not' and 'and'?

Male: You said earlier that I can do everything based on 'and' and 'not', and I can't.

Lecturer: Bill is going to tell us how to do that.

Bill: Sorry?

Lecturer: Because I can't think anymore.  $(P \text{ or } Q)$ , so either – well, actually, that's quite easy. It's ' $\text{not-}(P \text{ and } Q)$ '. ' $\text{Not-}(P \text{ and } Q)$ ', the truth conditions of that is branching, and on one side you've got ' $\text{not-}P$ ' and on the other side you've got ' $\text{not-}Q$ '.

Male: No, that's not right, because you may have neither  $P$  nor  $Q$ .

Lecturer: You're quite right. So all I need is ' $\text{not-}(\text{not-}P \text{ and } \text{not-}Q)$ '. Then I get  $(P \text{ and } Q)$ . Happy?

Male: Sometimes! (Laughter)

Lecturer: Ye of little faith. (Laughter)

Male: Why on earth would I have faith after the previous discussion?  
(Laughter)

Lecturer: I told you, you have to have to have faith, but it depends on what you have faith in! You must rely on the uniformity of nature.

Anymore on the Principle of the Uniformity of Nature or can I move on?

Male: I have a really general question, if no-one else has.

Male: Mine was a general question.

Lecturer: Well, let's quickly do the paradox of confirmation then.

I see a black raven, a black raven, a black raven. I see lots of black ravens. So I form the hypothesis 'All ravens are black.' Do you remember that? By inductive reasoning I generalise, from my observations, to a hypothesis like that.

Now I want to *test* the hypothesis. So I see another black raven, another black raven, another black raven, and so on. Each one of those sightings is confirmation.

But white gym shoes are every bit as much of a confirmation, logically speaking, of 'All ravens are black' as a black raven. That's because 'All ravens are black' is logically equivalent to 'All non-black things are non-ravens.' Do you see that that ['All ravens are black'] cannot be true unless that's ['All non-black things are non-ravens.'] true?

Male: Yes.

Lecturer: And that cannot be true unless that's true (vice versa). That's what 'logical equivalence' means.

So, given the logical equivalence between those two, if logically speaking the sighting of a black raven confirms that ['all ravens are black'], and the sighting of a white gym shoe confirms that ['all non-black things are non-ravens'], then the sighting of a white gym shoe must confirm that ['all ravens are black'].

Female: Sorry, how did you get from the first line to the second line, from 'All ravens are black', to, 'All non-white things are ravens'? Because the application of black in English up there means to the raven, not black ...

Lecturer: If all ravens are black. So here [drawing Venn diagram] are all black things, and all ravens are in there, so black things, those are ravens. All non-black things, so that's everything – no, I can't do that.

[inaudible discussion]

Female: It's all things. It's all things which are not black.

Male: You're right.

Lecturer: These are non-black things. So if all ravens are black, that means everything in the category of ravens is in the category of black things. This says that if anything is in the category of non-black things, like my shoes, it must be a non-raven.

Female: Okay. Yes, sorry.

Lecturer: How can my shoe, which isn't black, be a raven, if all ravens are black?

So all ravens are black and all non-black things are non-ravens. Those are true in exactly the same way. They are just notational variants on each other.

Therefore, if logic tells us that the sighting of a black raven confirms all ravens are black, logic should also tell us that the sighting of a white gym shoe confirms all ravens are black, doesn't it? But it would be completely mad to go out looking for ...

So what is it that tells us that the sighting of a white gym shoe *isn't* a confirmation of 'All ravens are black'?

Male: Well, strictly it is, but it's a very weak confirmation.

Lecturer: Well, so is the sighting of a black raven. That's a hugely weak confirmation of 'All ravens are black.'

Male: Not as weak as white gym shoes. (Laughter)

Lecturer: No, it is. It's exactly the same.

Female: It's not the same *now*.

Male: Not given the number of ravens there are ...

Lecturer: Ah, okay, no. Now you're talking. The fact is -

Male: That's why it's weak.

Lecturer: The sighting of a white gym shoe could confirm all ravens are black, if this were a world in which there are only ten things, and one of them is you, and you knew that – now, I've got to get this right.

There are only ten things. One is you. There are nine ravens and nine black things. That's a logically possible world. It's not an empirically possible world, but it's a logically possible world. If you saw a white gym shoe in that world, you would know quite conclusively that all ravens are black, wouldn't you?

Male: Wouldn't that mean that you're the white gym shoe?  
(Laughter)

Lecturer: Did I say that? Sorry! Let's not include you in that. But can you see where I'm going?

Male: Yes.

Lecturer: It's our belief that we live in a world in which there is more than ten things. There are more non-black things than there are black things.

The moral of this story is that whenever you're designing an experiment, whenever you're trying to test a hypothesis, you bring in all sorts of auxiliary assumptions, of which you are totally unaware, and when you're testing your hypothesis you are also testing all those auxiliary assumptions.

It would be wrong of you to think that if you get something that makes you think the hypothesis is wrong, that actually it might not be one of your auxiliary assumptions.

When I teach scientists this – when I'm teaching philosophy of science to science students – what I'm trying to do is to get them to think about the assumptions that they're making when they design their experiment, and to make them realise that when they're testing that hypothesis they're also testing all these other things.

As a matter of fact, they could be false, just as the hypothesis can be false.

Female: One of the things we talked about before was about is there not some sort of weighting you can give to something when you keep proving it right?

When you say, for example, 'We've had 100 experiments and gravity always takes place at this level', whatever, obviously you would run another experiment somewhere else at a slightly different level, but you could keep some sort of a record that's saying this hypothesis is based on ...

Lecturer: Well, every sighting of a black raven is slightly more confirmation of 'All ravens are black', but until you can show that there is no non-black raven you haven't proven the hypothesis 'All ravens are black.' The fact is, again bad example, there are albino ravens.

Female: But doesn't the fact that there have been in the world 100 instance of those ravens, or whatever it is, isn't it something that you're building on? Is it not increasing the likelihood if you've ...?

Lecturer: It's increasing the likelihood, but with induction you never get certainty, because you never know that around the corner there isn't a black swan.

Female: Is there some way though to reflect: 'Here is my hypothesis. A hundred experiments have been run on this and proved this one right. Here's this other radical one – we've only done two so far.'

Lecturer: Which suggest the hypothesis is wrong?

Female: No. It doesn't make any assumption about right or wrong, but it shows weighting.

Lecturer: But yes, I repeat, the more confirmation you have the more likely it is that your hypothesis is correct, but you never have *certainty* that the hypothesis is correct.

Female: Yes, I just mean ...

Lecturer: You can only be certain of the *incorrectness* of your hypothesis when you find a counterexample.

Female: But could you say, 'Hypothesis, a hundred incidences, hypothesis, two incidences'?

Male: Yes, there is some hierarchy.

Lecturer: Yes, of course you can. Of course you can. I've been saying that for the last three minutes! Agh! (Laughter)

Male: Well, there is such a hierarchy. The mathematicians have a hierarchy of theory, hypothesis ... There are a set of ...

[inaudible discussion]

Lecturer: 'All ravens are black', based on the sightings of lots of ravens. Now, more sightings of black ravens confirm this hypothesis, but it never shows that it's definitely true, because you might still come across a white raven. Well, we know that there are white ravens, so it's a bad example, but ...!

Let's have a general question.

Male: General questions?

Lecturer: (Laughter)

Male: I've got one – not a logic question. As a professional philosopher who does lots of things, e.g. ethics, I was just wondering how you find logic helps, and indeed maybe there are areas that you work in where it *doesn't* help.

Male: Sorry, that was going to be precisely my question as well, because in lecture one you described certain arguments that aren't susceptible to propositional logic. I was wondering is this really useful in the real world, if you have an ethical problem that you're grappling with.

Lecturer: I use it all the time.

Male: That wasn't my question. My question was in philosophy, not in the real world.

Lecturer: Yes. Well, we can talk about philosophy. What is in the real world? (Laughter)

Male: [inaudible].

Lecturer: When I'm thinking about ethics, I might come across a sentence that I don't quite understand. It's a very complicated sentence. I'm not sure where it's coming from.

I will construct a truth table for that sentence, in order to see what the truth conditions of the sentence are. 'Under what conditions is this sentence true or false?' Because until I know that, I'm not going to understand the sentence.

Some sentences in philosophy, as you will know, are very complicated, and I can use this to break it down.

This is baby logic. I went out of my way to tell you that you will *not* be able to go away from here and use it to say whether a leader in the paper is valid or hugely unlikely ...

Male: But if I was grappling with the problem of whether we should invade Syria, and I brought to it the ideas of a just war theory, and constructed an argument, would that be applicable to propositional logic?

Lecturer: It would be applicable, but you could *also* construct an argument for the opposite conclusion. There would be a use for doing that exercise, but you wouldn't be able to decide a question, because you would start to look at the premises.

Male: Yes, because then you could disagree with one of the premises, or agree with one of the premises.

Lecturer: Exactly. What it does is it enables you to clarify what you're thinking. You can use this to apply some rigour.

It's not just this. This is one form of baby logic. Then you've got the predicate calculus. You can't use the predicate calculus without using the propositional calculus. So you've learnt the propositional calculus: you could now go away and learn the predicate calculus and immediately the arguments that you can formalise will double, if not quadruple.

But if you want to do all of them, you would also have to do modal logical, deontic logic, probabilistic logic. There are so many other logics that you would have to learn, many of which we don't understand.

Male: But this is actually useful in the real world, under certain circumstances?

Lecturer: I use it all the time.

Male: Oh, okay.

Lecturer: Yes. If I am trying to understand the leader in a paper, I will formalise it as much as I can, and I will analyse it, and then I will try and formalise it ...

Male: Well, we did that last ...

Lecturer: It would give me an understanding. Even though I wouldn't be able to formalise it, almost certainly, I would be much, much clearer.

Lecturer: I would analyse it first.

Male: Oh, yes. Well, that was lecture one, wasn't it?

Lecturer: Yes. So, put it out logic-book style, because you can do that with inductive arguments, even though you then can't formalise it.

Female: When you say that many of these logics we do not understand, do you mean the lay person or do you mean the person that created the logic?

Lecturer: No, I mean philosophers. Look at this. 'Doing A is right.' Morally right, I mean. 'Therefore, I should do A.' Well, that actually looks like an entailment, doesn't it, a deductive entailment? It looks as if, if that's ['doing A is right'] true, that ['I should do A'] must be true.

But if it is an entailment, it's a very, very odd one, because it means that 'is morally right' is itself a truth-functor. It's a really interesting question. Kant says this.

Hume says that whenever we act, we act on a desire of our own. So, 'Doing A is right. *I want to do right.* Therefore, I should do A.'

Now, we can see that that's an entailment. That's a bit of practical reasoning that we would all accept. 'Doing A is right. I want to do the right thing. Therefore, I should do A.' That's a Humean bit of reasoning.

Kant comes along and says something very -

Male: You wrote Hume saying 'Kant'.

Lecturer: Yes. That's what Hume says.

Male: Oh, sorry.

Lecturer: Then Kant comes along and says, 'Well, hang on a second. You're implying here, if you think that that's necessary, that you might *not* want to do right, and what that shows you is you don't understand 'right' at all.'

Because if you really have the concept of morally right, then if you think doing A is right, your desires just fall out of the picture. You *should* do A if doing A is right, and what you *want* is just completely irrelevant.

Female: But thereby becomes a distinction in the actions. The former is an instinctive action and a reaction, but would you, because it's your disposition to ...

Lecturer: No, absolutely not. We're talking about moral reasoning here, and we know how difficult it is to get children to understand moral reasoning, because it's not an *instinct*.

The idea is that once you've understood the concept right, and a child probably doesn't understand it until it's about ten (in this country we say that the age of legal reasoning is...

Male: Ten.

Female: Criminal responsibility.

Male: The age of criminal responsibility.

Lecturer: Yes, but there's actually a ... Anyway, never mind. The age of criminal responsibility is the age at which you are deemed to understand the difference between right and wrong.

A child of three doesn't do this because Mummy won't like it, and a child of five doesn't pinch sweets from the sweet shop because she will get caught, or Mummy won't like it or something.

If you don't pinch my purse because you think you might be caught, there's something seriously wrong with your concept of *right*, because by this age you should be not pinching my purse because it's the wrong thing to do.

We know lots of people who have a concept of 'right' that actually isn't a moral concept, if they're still trapped in the 'I might go to prison', 'I might be discovered', or something.

But actually, the right concept of 'right' is one where your desires drop out of the picture, in which case we've got an entailment – a piece of practical reasoning – that goes from a belief straight to an imperative. I get the 'should' straight from a belief, without any desire at all.

Now, that's really, really interesting, and that's a problem in deontic logic. 'Deontic' means 'moral'. Deontic logic would say

that that is an entailment, and logicians would dearly like to know quite how that works, and indeed whether it does work.

Lots of people continue to think Hume is right. I really see that argument. I can also see why Kant wants to say, 'If you think that needs to be in there, you haven't yet understood what 'morally right' means.'

Male: I just think, those first two sentences there, it seems to me that the second sentence is what is meant by the first one.

Lecturer: Well, that's in a way of putting the fact that it's an entailment. Can you actually believe that something is right, morally right, without thinking you *should* do it?

It's certainly true that we can think something is right without *doing* it. We do that all the time. But that's different. The fact that you think you *should* do it will manifest itself in guilt if you don't do it.

So you still think this entailment goes through, it's just that you give in to temptation and you don't do A, even though you believe you should do A. So you feel guilty.

Male: You might think you should do it not because it's a moral law, but because you have internalised instructions.

Male: Oh, yes, ...

Lecturer: So you're not doing it because Mummy would be angry with you.

Male: Mummy is in your head.

Lecturer: Well, if Mummy is in your head, then you are doing it because Mummy will be angry with you.

Male: You're saying that could be the basis of morality?

Lecturer: Well, it *is* the basis of morality. I mean you've got to start by saying, 'No, darling. Lying is wrong.'

Male: Not as Kant said, 'There's a moral law.'

Lecturer: That's what Hume would say, and I've given the argument against that in saying what Kant would say.

Male: Well, I left out the middle bit as well, so you could straight to ...

Lecturer: No, the middle bit is implied in what you're suggesting. So you're internalising. In other words, you're not conscious of this, but it's still there, and that's exactly what Hume says. You don't actually have to be conscious of it.

Male: To go back to my original question, then, you can use propositional logic to construct an argument, not just to analyse it? Or you put forward an argument and then reanalyse it?

Lecturer: Well, let's do it now.

Male: I don't know what I think. I want to construct an argument to convince myself of something.

Male: Yes, I think ...

Lecturer: 'Allowing innocents to die is wrong. In Syria innocents are dying', etc. etc. I could construct a deductively valid argument very easily. It wouldn't be a huge amount of use to me, except that it would be an interesting exercise to do it. But I would then have to turn around, as a good philosopher, and construct ...

Male: And analyse it?

Lecturer: ... well, and then construct an argument for the other side.

Male: A counterexample.

Lecturer: When I'm listening to people on the radio talking, that's what I'm doing. I'm hearing an argument here for 'P', and an argument here for 'not-P'. Both the arguments might be good arguments. So then I'm thrown back to looking at the premises of the arguments and asking whether they are ...

Male: Yes, because in the end, you end up by questioning the premises.

Lecturer: I've taught you nothing about questioning premises. All I've taught you is how to identify a valid argument. You've got to look at premises on your own. That's what science is for.

Male: Is it?

Lecturer: Not always, but sometimes.

Male: Or ethics even.

Lecturer: Or ethics, yes.

Any other questions?

Are you happy with the paradoxes of confirmation?

Female: Yes, thank you.

Lecturer: You see why it's a paradox?

Female: Oh, yes.

Lecturer: I can give you another one if you like – grue – but we've only got half a minute left, and that would be optimistic!

If nothing else, the bar is open. (Applause)