

Trust and Deception

Title: Conceptualising Trust and Deception (6 October 2015 Part 1)

Date: 06 October 2015 (09:00-12:30)

Duration: 55 mins

Speakers:

- S1 - Guido Möllering
- S2 - Nicholas Wheeler
- S3 - Nicholas Wright
- S4 - Geoffrey Bird

S1 So I'm going to try and say a few general things about the relationship between trust and deception, and suggesting a rather peculiar relationship because I'm claiming that trust and deception may enable and prevent each other. So we have to go fast, I think – try to be fast. How much time do I actually have? 15/20?

S2 I mean, ideally, if you could do no more than 15. Ideally.

S1 OK, well, so I have to be really fast. So I started thinking about trust ages ago and about deception around eight years ago, when I got a chance to be involved in a book project – a really nice book project on deception. And we had meetings at the Santa Fe Institute and really with the deception researchers and that was, for me, very new and very exciting. And so unfortunately the ideas I'm going to present today are mostly the ones I developed back then, but I keep on thinking about how trust and deception are related, and I think it really helps us to understand trust better if we connect it with deception.

Trust and Deception

Anyway, so the ideas go back quite a long time, but I'm trying to suggest some new ideas. I have, like, five main points that I want to make, or that I am making as a basis for this, but I'm not going to read them now. I'm just trying to summarise them in my own words in a moment, or [0:01:17] but in a different way, and then maybe we come back to the slide and, you know, if we need to be very precise, I will need this slide.

So the first thing is about how people may recognise trustworthy people, or rather untrustworthy people, or differentiate between people who are genuinely trustworthy and not genuinely trustworthy, and this is the question of which signals of trustworthiness are actually reliable? And this is all not very sharp but I hope you can see that the – here's the image of the book *Streetwise: How Taxi Drivers Establish Their Customers' Trustworthiness*, by Diego Gambetta and Heather Hamill, and they – Diego Gambetta, you know, is a big contributor to trust research going back through the 1980s. He got interested in the signalling theory of trust, and together with Michael Bacharach wrote a theoretical paper on reliable signals related to trust, but then with Heather Hamill he also did ethnographic work with taxi drivers in Belfast and New York about – you know, asking them/doing field work with them – how they actually know whether a passenger – a customer is somebody they should take or not – or, rather, not take. And so the whole – the core idea here is that some signals – well, 'cause signals of trustworthiness – then you may think of a lot of signals of trustworthiness but the crucial question here is which of them can be faked and which of them cannot be faked?

And the theory – the signalling theory as they apply it here is that trust should rely on those signals that are reliable because they are hard to fake, or it would at least be very, very costly to fake, and so they are the reliable ones, and the ones that are easy to fake are not reliable signals. And so the examples they would use, or an example that you

Trust and Deception

could use is that in general – for example if, generally speaking, women are more trustworthy than men, then it's relatively – relatively – hard for a man to fake being a woman just to be trusted. Right? You can do it. You can drag or whatever, but it's a very big effort. Or, for example, older people are often seen as more trustworthy than younger ones and, you know, your age is also very hard to fake. Not impossible. You can go some way, but – and the whole theory revolves around the idea that the taxi drivers have a feeling for which reliable signals they can use, and they actually do use them deciding, for example, whether to take somebody or not.

It leads to some interesting stories like in Belfast on a Saturday night, if you're not drunk, you're suspicious, right? 'Cause it's – or in New York, like, 14-year-old teenage girls are the most untrustworthy customers you can get because they very easily jump out of the car without paying and so there is – but the taxi drivers develop this kind of idea which passengers are trustworthy and which ones are not.

This is all fine but for me this is far too rationalistic, because in real life the signals that people are sending and receiving are very messy and you – you know, when you are in Birmingham on a Saturday night, you get – you know, it's chaos, right? Sometimes.

[Laughter]

You get people who are very drunk but maybe not necessarily very aggressive. Maybe they are really even helpless and so on and they say stupid things but it's really – I mean, you should either say the taxi driver should basically not be at work that night – should just stay at home – but if they are working they have to figure out amongst some very messy signals what to do and when really to say no to a passenger and when to kind of work through this mess of what the customer is offering in terms of signals.

Trust and Deception

And then of course the other way around. You've just got your own taxi scandal here. Maybe it wasn't that prominent but I found it on the web. Birmingham taxi driver scandal. This is from last week – two weeks ago or something like that: '114 cabbies with criminal pasts granted council licences'. So it goes the other way, of course, as well. Passengers are wondering how trustworthy the driver is and are picking up all kinds of signals and one really important signal that should be hard to fake is the licence that they have. The little card with the photo. And that, if the city council isn't really paying enough attention to whom they are giving the licences to, then what is that signal worth if it's still kind of not really proper – maybe?

OK, the bottom line is a lot of effort is – or a lot of thinking is also invested in how we – in thinking about trustworthiness and how trustworthiness and recognising trustworthiness is related to deception, but for me that is only a very small part of the story, and because it is much more – it is not much – less like somebody's standing there with a checklist, seeing if they can be trusting somebody, but it's about the relationships and it's about what happens between the trustor and the trustee. And some colleagues in Berlin – one PhD student there just submitted a PhD thesis on trust with taxi drivers, and argues very much that this is a really very rough – a very rough approach to explaining what happens in the taxi, because really the taxi driver and the passenger often negotiate while are they already [on driving - 0:06:50] how much trust they can have and how much they trust each other and what it is all about – what the relationship is all about. So I'm going to argue that we shouldn't look at this as such a static thing of just recognising fake signals, but much more a relational thing of how the – both sides of the relationship work on the relationship and work on how much trust there is and also work on where the deception occurs or doesn't.

Trust and Deception

OK, so the title says, 'Trust and Deception Enable and Prevent Each Other'. Very quickly: deception destroys trust. OK? I don't need to talk long about that. When we're deceived, we lose trust in those who have deceived us. Trust opens the door to deception, and that's also mainly easy. You know, when we trust, we lower the guard, and that makes it, in principle, easier for the other side to betray us/to deceive us, maybe even without us noticing it.

OK, so these are easy? Are they OK? Do you need more explanation of those two? The next ones are less obvious. How trust reduces deception or prevents deception. Now it's getting interesting. Now it's getting to the point of what the relationship is all about and how people perceive the relationship and how valuable the relationship as such is. And so if the relationship is seen as valuable by both sides, then trust should reduce the concurrence of deception because people will not deceive the other side because they don't want to lose their trust. You know? If this is easy, then it's also easy to imagine that people will not deceive because they're interested in keeping the relationship and keeping the trust. So, actually, by trusting Nick, I can kind of make him trustworthy in the sense that he then is going to try not to deceive me 'cause he doesn't want to lose that trust. But it's, of course – the big question is whether he's interested in having that trust. But if there is a relationship that is valuable, then the higher level of trust that is in the relationship should make it less likely that people will deceive each other in that relationship.

And now maybe the weirdest one is that deception increases trust. Because we just said that it kills trust, didn't we? Now, this point here – and I'm also going to send this paper around, which is like the longer version of the chapter in the book, and this is a bit more academic and you can download it for free. So the point here is that trust involves – always involves deception and self-deception, and so, when people find out how much they can trust each other, they – yes, they present

Trust and Deception

themselves as trustworthy, and to some degree this is always just a partial representation of what people really are. So if I want to find out that I can trust you, then I'm hoping, actually, they you are going to make it easy for me by showing me your trustworthiness and displaying only those signals to me that will make me believe that you're trustworthy. So you are keeping a lot of things untold and you can say that maybe that is actually deceptive, because I should know everything about you before I decide to trust you. No, you are deceiving me in a way that you are presenting yourself in a favourable light. And so will I, and that can be this kind of fiction of us both being trustworthy, which is partly a deception or self-deception, can help us to grow the trust. So it's this idea – there is this idea that people want to trust and so they want to have a favourable image of the other and the other one will help them to do it. Of course, con-artists do exactly that, so it's risk, of course. If you think about it in this way, con-artists, their main skill is to make us trust first before they deceive us and before they betray us. But in other contexts we have to help each other to trust and we have to make leaps of faith and we have to work with incomplete information and biased information, even. We do – we have a confirmatory bias in all of these things that we can find in psychology, but to some degree, the fact that trust is an illusion or rests on illusion, as Lumen would say, makes it possible to start with. OK? So that, in a nutshell, is this...not too complex but rather interesting relationship between trust and deception.

But how – so how can we then find out – or how can we go a step further and find out how it really then works out? And I think I'm probably almost out of time now, but....

S2 No, you're OK.

S1 But this part is now important and is also a little bit more recent thinking rather than what's already in those applications. Now, we have to

Trust and Deception

consider the conditions under which trust matters. And maybe – I don't know which definitions of trust you're using, but to me one of the things that always comes up is that trust matters when there is vulnerability and uncertainty. Yeah? So it's because in a relationship we have vulnerability/uncertainty that we require trust as a way of dealing with that. At the same time these conditions make deception easier, so if it's easier to deceive people who are vulnerable and who have, you know, limited control over the actions of others, then this explains the kind of negative relationship, too, but it's important to see that deception requires trust – it requires vulnerability and uncertainty. If we have certainty, we cannot be deceived. You know what I mean? If there is certainty – if we knew everything – if we had perfect knowledge about others, the others wouldn't be able to deceive us. It's only because of the condition of uncertainty that there is a possibility of deception. Yeah? And if we were invulnerable, then the deception wouldn't matter, either, because then they can deceive us as much as they like but it will not hurt. So the same conditions that we need to talk meaningfully about trust are also required for deception. And so I think that connects the two concepts, or should connect the two concepts, and be an interesting idea, hopefully.

And on the other hand, I would like to suggest that we need to think about the norms of solidarity and responsibility. Mutual responsibility and solidarity. So if in a relationship there is a norm that you are responsible for each other and you will support each other, then this norm supports the trust that you have in the relationship. And this norm also...makes its – explains why people will not use deception, because they want to live up to that norm. Of course, this is very, very – in international relations, I'm aware, this is the big question. Can we assume that there is neutral responsibility and solidarity between the heads of state or even the entire nations or states that they are leading? OK? So this is – and I just want to highlight that for now and see what you say about it, but I want to suggest one more thing, and

Trust and Deception

that is the connection between the conditions and the norm. And this is now where I'm going to get a little bit – maybe too romantic, but I think that, if the actors involved in a relationship realise that they both experience vulnerability and uncertainty, then this could give rise to solidarity and to mutual responsibility. We have this in business relationships. Very mundane business relationships – buyer and suppliers. I get told by these managers involved in that – they say like, 'We are trying not to make our lives unnecessarily difficult for each other. You know, we both know that, you know, we cannot control each other perfectly. We know this, and so we have made kind of implicit agreement that we are not going to make it difficult. We are just going to help each other where we can, look after each other, just to make the relationship smooth and easy and better and...', well, maybe in business that doesn't mean that they love each other, but they are feeling certain responsibility and solidarity against the background of this uncertainty and vulnerability.

So: 'Whether trust and deception enter into a positive or negative interaction depends on the level of mutual responsibility and solidarity of the relationship. The willingness to accept mutual responsibility and solidarity may be supported by the recognition of mutual vulnerability and uncertainty, in which – but in which—' and this is more than encapsulated interest, according to Hardin. It's more than mutual hostage-taking. It's a more general condition that we're talking about here. '...in which context and at what – but, in which context and at what levels is it reasonable to assume that all sides are willing to commit to a positive relationship?' And, again, what if they are only faking this solidarity? So the problem doesn't go away just by stating this kind of possibility. There's still a deception there and there's still a need for a leap of faith there. Yeah, and we can consider whether it's an inter-personal thing amongst these leaders or it's an international thing. The flags are behind them; they represent various countries.

Trust and Deception

OK. Now, just to make it – one minute.

S2 OK.

S1 To make it even – I don't know whether this is all for you kind of familiar stuff and the way you think about deception anyway, or whether this is rather unusual, but I want to – I already said that this idea here is a little bit romantic. Yeah? So we have to – vulnerability and so we look after each other – yeah. But I want to make that even stronger with this sort of Shakespeare sonnet, which it's just lovely and it's actually also in the red book there.

So what – so how can there be a positive relationship between trust and deception?

When my love swears that she is made of truth
I do believe her, though I know she lies,
That she might think me some untutor'd youth,
Unlearned in the world's false subtleties.
Thus vainly thinking that she thinks me young,
Although she knows my days are past the best.
Simply I credit her false speaking tongue:
On both sides, thus is simple truth suppressed.
But wherefore says she not she is unjust?
And wherefore say not I that I am old?
O, love's best habit is in seeming trust,
And age in love loves not to have years told:
Therefore I will lie with her and she with me,
And in our fault by lies we flatter'd be.

So this is a happy, deceptive, entrusting, love relationship. And, finally, another person who studied the trust of taxi drivers and passengers, Jim Henslin, actually uses this definition of trust, which is really

Trust and Deception

interesting, I think. This is how Henslin defines trust: 'When an actor has offered a definition of himself and the audience is willing to interact with the actor on the basis of that definition, we are saying trust exists.' So, basically, trust is going beyond or simply crashing through the possibility of deception and false impressions and it's basically saying, 'OK, I take you to be what you are showing me that you are, and we work on that basis.' And therefore Henslin's definition of trust, which includes very strongly (to me, at least) this idea this may be fake; this may be a fiction; this may be even deceptive. Thank you for your attention.

S2 Thank you, Guido.

[Applause]

OK, our second speaker is Nicholas Wright. Nick, as he said, is a senior research fellow in the ICCS. I'll also do a quick plug for him. He'll be talking next Tuesday at 12 o'clock in Room 417 back in Muirhead Tower on 'Knowing How the Other Thinks: the Brain and Influence in International Confrontations', so do come along to that. But today his topic is 'Organic Deception and its Limits: The Human Brain, The United States and China'. Intriguing title. Nick.

S3 OK. Thank you very much. So, originally I was actually going to give this talk with Nick, and then two weeks ago I discovered I was giving this talk on my own, so it will involve slightly more questions, perhaps, than answers, but.... So today I'm going to talk about deception. So deception is wide-spread both at individual level and at the state level, and we need to think about both of those levels because the fact that something matters at the individual doesn't mean it matters at the state level and vice versa. And I have a practical bent, so I'm also at the Carnegie Endowment for International Peace, which is a think-tank in Washington, and I'm very interested in practical, real-world matters, and so what I'm interested in here is not only just to say that deception is wide-spread at the individual and state levels but also to show what

Trust and Deception

limits deception and what shapes deception, because that's, in many ways, the most important thing.

And I'm going to quickly go through some different ways of thinking about deception at the individual level and at the state level in terms of different aspects of decision-making and a biological, psychological and social sort of take on decision-making.

And so we can think about humans as organic decision-makers. So an organic account of decision-making in individuals and in states can be thought of – so the account's organic but is encapsulated by the word 'organic' in two ways. So basically because organic means 'relating to or derived from living matter'/'characterised by gradual or natural development', so those are the two ways in which I use the word 'organic', and we can think about organic individuals and organic states. So organic individuals – so the first thing to say is that, you know, we are organic in that we have human brains and bodies and human choice arises from the brain, and so we can think about deception and we can think about sort of psychological experiments that looked at deception.

So a lot of the best of these have been done by Dan Ariely, and a chap at Harvard, as well, and one of the key things that they've done is they're basically showing the people their costs to deceive. So essentially they do experiments where, for example, they have to – these are simple, laboratory experiments where they will have to say – do a simple task and they have to be honest about how many times they succeed – how many general knowledge questions they got right or something along those lines. They can lie about it. They think the other people don't know that they've lied or the experimenter doesn't know that they've lied, and basically what happens is that they lie, and then subsequently when they're told, 'OK, well, now you need to accurately predict how good you are at answering general knowledge

Trust and Deception

questions', they're influenced by their own lie and they end up bearing a cost because they're influenced by their own lie.

So there's a lot of that interesting research, but basically people lie a lot and people pay – you know, pay costs [essentially, because - 0:22:52] of that lie.

We can also think about it in terms of behavioural game theory, so this is not game theory as such; this is a psychological take on how people are involved in strategic interactions, and so basically, for example, as Guido was talking about with screening – so, you know, we don't have to invent all sorts of new ideas. These ideas have been around for ages, so there's screening, for example, which is, how do I know? If I'm interviewing a student, for example, how do I know if they're a good student or not a good student? I can think about it by, you know, asking them to do costly things, for example, and good students tend to do more costly things, for example. So we can [0:23:28] screening, and we know if you then get people in experiments to – and we test them in these types of environments – these screening environments, for example, then we see that human beings do actually basically use – they play – in many ways, as expected, people lie about their type and people try and understand what their type is.

And just to say that people often lie about things because they're things worth lying about. So a great example from Washington is David Petraeus. So what did he lie about? So he was the head of the CIA. This is an enormously successful guy. He was behind the surge in Iraq and then he's the head of the CIA, and what does he lie about? He lied about sex, and there was this great image of his – unfortunately; I hate to say it – very – in the tabloid press they had his very frumpy-looking wife and then this sort of very attractive woman who he had his affair with. And so why do people lie? Often they lie about things that really matter – things like sex, or they lie, you know, on behalf of their

Trust and Deception

family. OK, so, based in our psychology and biology, we can think about that this is just something that humans do.

The second thing we can think about is that an individual develops organically. So children start to lie – and it's an important thing in child development, that children start to lie. I have a two-and-a-bit-year-old daughter, and she's just starting to understand that she can deceive me. She's awful at it at the moment, which is a good thing, but basically she's just learning that she can deceive me. And this is just an important thing, and all kids have to learn how to lie. And then they have to learn how to control their lies – when to lie, how to lie better and so on and so forth. And there's also – this is a great opportunity for future work, is thinking about how does lying change over the lifespan? And there's very, very little on that at the moment.

We can also think about – so an individual develops organically, and it's not just about human nature; it's also about nurture, and so we can think about how does culture affect lying? And so, again, this something we just don't know very much about, but.... So Guido comes from Germany. I don't know – do you come from East Germany or West Germany?

S1 West.

S3 So you come from West Germany. So you are likely to lie less than East Germans, apparently. So, for example, you can get people to roll a dice – 'cause they did this in Germany – [0:25:45] experiments were done [0:25:49]. You get people to roll a die and then they have to say, you know, what they were going to – you know, well, I can't remember what it was. It was like you have to say whether you got a six or a one or whatever. And, basically, East Germans were more likely to suspiciously get loads of really great rolls. They got a lot more of those than West Germans, for example. And the idea that they were putting

Trust and Deception

forward was that because – these were people who grew up in East Germany with East-German parents, relative to those who grew up in West Germany with West-German parents, and the idea is that those who grow up in these types of societies where you have to lie all the time, you have an inner self and an outer self – and this is like, you know, me, when I speak to my friends at Peking University, you know, they say obviously people in China will lie more because you have to lie. You have an inner self and you have an outer self. And so, anyway, the idea is the culture affects the amount of lying. But, again, we just don't know very much about it, and that's actually a hugely important thing.

So, for example – I mean, we don't have to go into it, but here on the right we have a slightly different thing. And I'm so sorry, Theresa. I always mention this terrible study about Greece, so I apologise profusely 'cause Theresa's from Greece. We have a great study—

[Theresa: I can take it. - 0:26:56]

S3 —which basically shows – it's just looking at something called a Public Goods Game, and when they looked only at northern Europe, and they also looked at China; basically what they showed was that – and this was a science paper that came out probably a few years ago (and this was before the financial crisis) – they basically showed that, as expected – they'd known this for some time from other science papers and whatever and everyone was very – thought this was a very solid finding – that people punish free-riders in games. But then what they discovered is that when they went to various places, particularly in the Middle East and in Greece, people will punish people who contribute a lot more. They will pay to punish people who contribute to the common pot. People who contribute more than average, those people are punished. Now, why on earth is that the case? Nobody knows. It's basically an empirical finding. We need an explanation for that. We

Trust and Deception

don't have a good explanation for it yet, but that's the kind of thing I suspect we're going to find if we look at lying in an interesting way across cultures.

I don't think we're going to have a huge amount of time here, but then – so let's think about organic group societies in the international system. So the fact that individuals lie doesn't necessarily mean anything about groups. OK? It doesn't necessarily mean anything about states. Because, you know, this is what people always say about – if you look at things like loss aversion – any of these things we show in the laboratory, or even if we show in the real world at the individual level, it could all wash out at the aggregate level. We have no idea whether it matters at the level of states.

So, what can we say? So do individuals matter? So I think that the [0:28:42] one example is individual choices of leaders do matter, and we know that presidents lie. There was a whole book all about presidents lying, and that was a great start of this book where it says, you know, there was this advert – I forget; you guys may know this better than me. I wasn't living in America at that time, so I didn't see – and, in fact, I only watched internet [0:29:02] lived in America, so I never saw any adverts, really, except from Amazon. But, basically, in – what they showed is – I think it was like: 'Watch my lips: no new taxes' and then it was – what was the next one? I can't remember what it was. And then it was like Clinton saying, 'I did not have sexual relations with that woman', and you know, it's like, basically everybody lies, was the gist of the thing. When you're a president, [0:29:24] FDI lied, and in fact one of the things I was going to do for this but I didn't have time is to go through every president and say some important way in which they lied. They don't even talk about Nixon in that book, really. There's not even a chapter on Nixon, because it's bloody obvious that he lied. That's how widespread it is. And if you believe that leaders matter, and I do believe that leaders matter, and even if

Trust and Deception

you just believe that if you, you know, transfer it – individual leaders don't matter; that just leadership roles matter, then the fact that basically all leaders lie and have to lie, and I don't even think that's a bad thing – then lying will occur at the level of international system and obviously we – it's not an American thing so, you know, 1956, you know, Britain and France – I don't know if you're familiar with the Suez, but basically we needed to engineer a reason to invade Egypt. So how did we do that? So the way we did it is we got the Israelis to invade and then we – then Britain and France sent a huge taskforce over to say, 'Now all we're doing is just being peacekeepers.' But we'd obviously colluded with them. And this is all very, very well known. We concluded with the Israelis. We basically told them to invade. So, you know, this is a British thing, as well.

And then, for example, think about China. So I'm going to give China a little bit more detail, but a key thing about China is that what's left of – so Chairman Mao is a critical person obviously in Chinese sort of post-1949 history, and he's critical to the Communist Party. So what is left of Mao's reputation? So nobody really believes that Communist internal reforms and way of organising the economy is a good thing. Nobody really believes that. Nobody believes the culture revolution was a good thing anymore. Nobody believes that the Great Leap Forward was a good thing anymore, really. Obviously there will always be some people who do, but not many. So what's left? The thing that's left for him is his foreign policy stuff. And, essentially – like, one of his key things was to invade North Korea to repel the Americans, when the Americans were getting close to the Chinese border of the Yellow River. And basically what we now know is that, like, Mao and Stalin were both very, very happy for the north to invade, and basically told the north to invade the south and allowed it to happen. And so the whole idea for the – that Mao was some great statesman for repelling the United States in the Korean War, which cost, you know, hundreds and thousands of Chinese lives, is basically based on the idea that this

Trust and Deception

war came out of nowhere and they were repelling the United States. That's not true. He was intimately involved in creating this war. So there's a great book by Thomas Christiansen, for example, that talks about that [0:32:08].

Another key thing, and this relates, for example, to work I've done on ISIS and so on, is that society is inter-generational. So you have societies organically developing humans, and the problem is that we have all these kids who are developing how to lie and how to lie well and then you get to university and you learn to do all sorts of clever, sophistic— like, you get to be a postgraduate, you get to – you want to do psychology, for example; you have to learn how to respond to reviewer's comments. You know, in a clever way that's not actually lying but, I mean, it is basically lying, isn't it? It's – sophistry is the Greek term. And – which always sounds a bit poncey [0:32:44] in Greek. I'm so sorry, Theresa, but it's true. And so – you know, we have to learn how to lie. And the problem is that even if we all came to the conclusion that lying wasn't a good thing or whatever, you will always have new generations of people coming up who will lie. And that's just, you know, the way of the world. And also you will always have mutation and adaptation, so you will always have new people who will come in who will lie.

OK. And just to make one very quick point about the challenge we face both as individuals and as – in fact, I'm not going to make that point.

OK, so think about the States now. So I'll give you a quick example of the US. So surprise, for example – the whole idea about surprise is one of the main ways you generate surprise is that you deceive somebody. So in the first Gulf War, for example, the US tried to deceive Saddam that they were going to invade via the sea, and in fact they didn't; they pommelled them by land. OK? And obviously Britain

Trust and Deception

did that in D-Day – tried to persuade the Germans they we were going to land on one thing and we landed on another thing. And that's basically central to the way that both the United States and China think about interacting with others in military confrontations. Indeed, in crises.

And thinking about China, so they explicitly – they have a much more psychology-centred approach to manipulating others and influencing others than, for example, the Unites States, so they, for example, talk about the three warfares (and this was endorsed by the CCP – Chinese Communist Party Central Committee – in 2003), and that involves psychological warfare, media warfare and lawfare. And so they explicitly want to manipulate other people's psychology, [0:34:40] their decision-making, and a key thing in that will be deception. And the reason for that is that they are very interested in how the weaker can overcome the stronger. They know that the United States has military superiority in a – almost across almost every spectrum – across the full spectrum, so the question is, how can the weak overcome the strong? The way you do that is basically with deception and surprise.

And of course you can think about classic texts, as well, so deception is essential a key component of [0:35:09].

And just to say – just to illustrate – so this is the wrong way around 'cause it was sent by a collaborator in China – or the right way round, depending – so what you see here is this – for example, this is essentially – so we did a look at all the databases and basically the Chinese talk a lot about psychological warfare, and there's been a massive growth of that recently. This is across lots of different databases that they use. This is academic journals, [BCs - 0:35:36], newspapers and so on.

S2 Couple more minutes, Nick.

Trust and Deception

S3 I'm hopefully going to finish in one minute. So what I'd say is that it's realistic that human individuals and states will essentially always lie and that I think what we have to acknowledge is that this – and this is like a key thing. We always get bogged down on this – is – so do you know about, like, David Hume and the idea of is and ought? So this – so he basically says that – when he's talking about ethics, he says, 'We always go – reading these things, I always go from a statement about how the world [is to - 0:36:10] a statement about how the world ought to be.' But we don't want that. What we want to do is we're saying, 'This is a statement about how the world is.' Now, there's an interesting discussion to be had about how the world ought to be – about the morality of it – but that's separate to how the world is, and so what I'm saying is that states and individuals lie; that is how the world is; I think that's how the world is going to be; and that's not how it ought to be, for moral reasons; and that's how it would be according to some rational-choice theory or other particular type of model, for example. This just is the way it is, and the way that humans lie and the way that states lie is determined by a number of factors to do with their organic growth and evolution and also the fact that they [0:36:47] human beings.

So we have evidence that individuals lie even when it's costly; deception is inherent to key aspects of strategic decision-making and there are lots of easy cases that you can look at; for example, what they love to call in America 'MILDEC', or like military deception, for example. And, you know, we can go on from there.

And I think the key question, really, is actually not whether individuals and states lie. It's, what limits and shapes deception? And, as I said, there are interesting ethical questions but I think we need to separate those off and think about those separately. For example, if you're a utilitarian, you basically almost – you almost certainly will believe in

Trust and Deception

lying. And what we need to think about is, though, that there are limits to lying. So culture shapes and limits lying, norms are important, and that matters at the individual level and at the state level, and these can be – lying can be used as a tool for state but you have to be careful about how you're using it. So obviously, you know, that's a part of the type of work that I do with the Pentagon and so on, is thinking about how to use these things as tools, but you have to be careful about how you use lying as a tool.

Thank you very much.

S2 Thanks, Nick.

[Applause]

Our third speaker is Geoff from the Developmental Psychiatry Centre at King's College, London, and Geoff's going to talk on how good are we at detecting deception? So that follows on quite nicely.

S4 OK, so I haven't got very long and there's a lot to talk about this topic, so I'm going to speak very quickly. Feel free to interrupt and ask questions if you want to. I'm going to talk mostly about our work, mainly 'cause I had the slides, and it's easier to describe quickly.

There's a very complicated answer to this question – how good are we at working out when someone's lying to us? So there's also a simple answer. The simple answer is that we're absolutely useless. So if I said to you for breakfast this morning I had toast, you've got no chance whatsoever of knowing whether I'm lying or not. It's just literally impossible.

There's a magic figure in this literature. So if you use these type of paradigms where I say something to you like 'I had toast for breakfast' or 'I'm very much a believer in Communism as an ideal way of governing societies', and I give you lots and lots of those statements

Trust and Deception

and I test lots and lots of people, then the mean percentage accuracy will be 54%. It's astonishing how consistent that is. It's absolutely astonishing. It's the most consistent thing you're ever going to see in psychology or any other thing that I know. It's really weird. So we must have done about – I don't know – 15 studies of deception now; only about four or five are published, I think, but anyway, we always find 54%. And everybody else does.

So why, is the question. And the answer seems to be it's not a product of the lie detector. Mm?

[Student: It's a coin toss. - 0:39:41]

S4 Yeah, so it's a coin toss. So 50% of a chance. And it doesn't seem to be a product of the detector. It seems to be that the reason why this 54%, so just slightly better than chance, is that there are some people that are hopeless at lying, and it's those that get caught out by most people. And actually, what I think is maybe that there's not some people that are hopeless at lying; even the people that are really bad at lying only tell a few lies that people can detect.

So the idea that you're going to be caught for lying from cues in your face or anything else is complete and utter rubbish. There are no reliable cues to lying. The only thing that's a decent cue is response latency, so if somebody asks you a question, how quickly you are to respond. Very noisy. Nowhere near hugely 100% predictive.

What about technology? Polygraph – 85% accuracy if you take a very kind view of the literature. Huge base-rate problem. So people use this for employment screening, when base-rates of lies are very, very low; it means you get just stunning amounts of false positives. The two big spies in the FBI – FBI one and the CIA – Richard Ames and Hanssen, both passed numerous polygraph tests.

Trust and Deception

FMRI – use of FMRI to detect lies – fantastic on the individual level. So if you give – you use machine-learning techniques; you ask somebody to lie in the scanner/tell the truth in the scanner, you put the data into a machine learning – any type of machine learning you like, you can then classify later lies with 98% accuracy. Soon as that person comes back a week later and different things in their brain changes/the humidity of the room changes, those models don't work anymore. Those models don't work for anybody else. You can break FMRI's lie-detection tool very easily. You just move around a little bit.

Nothing works, basically. We're never going to be able to detect someone's lies with a decent degree of accuracy. People are now looking at different kind of interviewing techniques. They're slightly good at getting people to admit that they're lying. They're not really going to catch liars, really.

One thing that's massively missing from the literature is a focus on liars. We're very much looking at lie detectors, not liars. And I think that's really interesting 'cause it seems that most of this figure – most of the variants in whether someone's going to get caught as a liar or not depends on the liar rather than the lie detector.

So we've done quite a lot of this and we've used this Deceptive Interaction Task simply because we can call it the DECIT task, which is quite cool. My PhD student, Gordon, used to be in advertising, so he's very good at doing these things. This is a false opinion paradigm. We get lots of people to come in. We ask them their views on kind of controversial subjects – euthanasia, abortion, pornography – before they know they're going to do a lying task, and then we put them in groups and we ask them to give their true or false opinion. So there's one sender and there's other judges around here, and so the sender either gives their true opinion or their false opinion and everybody else

Trust and Deception

makes a judgment, basically. We give [0:43:04] prizes for the best liar/the best lie detector, to try and motivate people. They're motivated anyway. This is a really fun task. It's the only experiment where we wouldn't really need to pay people. Everything else is stunningly boring in our lab.

We've come up (I think we were the first) with a cool way of analysing these data. So signal perception theory is a nice way of analysing – this has been done before – in the receiver role. So what this enables you to separate is people's ability to detect perception, so to separate out lies from truths, so this is just everybody's [0:43:41] detection, and their bias, so this is the calculation of the bias. This, in essence, is how likely they are to say truth or lies irrespective of whether they actually heard anything. So if somebody's very, very gullible or very, very trusting, they will say truth all the time. If somebody is very, very non-trusting or paranoid they will say liar all the time. And that's completely independent of their ability to discriminate these two types of statements.

What we can do, which is quite cool, I think, is invert this for the sender role – for the liar role. So again we get two parameters. We get this d' parameter – what they call d' Sender – this bias measure. The bias is very interesting. So the bias is, before you open your mouth, really, or ignoring what you're going to say, how much do people believe you? Your credibility. There is amazing individual differences in this. To some poor individuals, before they've even said anything they're judged to be liars. No matter what they say. Other people have like a halo. They're just assumed to be wonderfully honest. Then you've got the [d' Sender] measure, which is a really nice, objective way of working out how good you are at lying. In essence, how much more – how much less credible do you appear when you're lying versus when you're telling the truth?

Trust and Deception

So these, I think, are really, really useful indices because they enable you to get an actual ability of lying, which we haven't seen before. So one of the things that we've shown is that people who are very good at lying are very good at lie detecting. So there seems to be a deception general ability. We've replicated this about eight times now and we always find it. I really want to stress that the people who are very good at detecting lies are crap at detecting lies, so nobody's very good at this, but in the variant in the population, the people who are better lie detectors are the ones that are better able to lie. There seems to be some kind of model that these people understand that they can apply to others and use themselves. It's a moderate/modest correlation – about point-three/point-four.

So there seems to be some deception general ability and, like I said before, the thing that seems to really matter when it comes to truth and lies is the response latency. So in our paradigm we say to somebody, 'OK, you were asked to talk about, say, abortion', and then we give them a cue – [snaps fingers] truth or lie – and everybody knows that cue's been delivered, then, and if you're lying you're slightly later to start and you speak for a shorter amount of time. And you can calculate this response latency difference between truth or lies and it seems to be that the people that are very good at lying, without explicitly knowing that's what they're doing, they're better able to reduce that gap between truth or lies.

Recording: Hi, my name's Amy. I'm a PhD student; currently single. I love animals and going to the theatre and on weekends I like nice walks in the park when the weather's nice. If it's warm, I like [0:46:42].

S4 One thing that good liars are good at is working out when to lie – who to lie to. So would you lie to this person? Put your hand up if you would. Excellent. OK. You're all rubbish at this. So she is incredibly gullible. She can't detect lies at all. So what we can show is that

Trust and Deception

people who are very good liars – again, it's a kind of modest correlation, but they are very good at working out the people that they should lie to. So we show them those videos – just a brief introduction to the kind of person – and so, 'OK, how good do you think this person is at lie detection?' So good liars are good at lying and good at working out who to lie to, which is, I think, two interesting things about liars.

How do you know whether someone's going to be good at this deception general ability? Whether it's telling lies or detecting lies? One thing I think is incredibly interesting is it's never IQ. We've tested thousands and thousands of people now. At least 20,000, probably more. It's not IQ. It's not emotional intelligence, however you want to define that, or measuring. Theory of [mind disability - 0:47:50] to represent other people's mental states: I think we've got really bad tests of this, but anyway, theory of mind predicts your ability to detect lies, not so much ability to tell lies. But, again, I think that's 'cause we can't really test theory of mind very well.

There are loads of stuff in the literature about personality traits that should correlate with lying. Psychopathy's the big one. Machiavellianism. The kind of use of lies to achieve your aims. Narcissism. Lie acceptability. Self-deception's very interesting. There are evolutionary theories of lying. Trivers is the main person in this area. Trivers argues that our ability to perceive ourselves evolved to make us good liars. So the idea is that people can't detect you're lying if you don't think you're lying, and so self-deception – there's huge amounts of stuff on this – very high-profile papers suggesting that self-deception helps you lie.

If we actually look at these we can now actually test people's ability and see what predicts what. Machiavellianism doesn't predict anything – lie-detection or lie production; narcissism doesn't predict anything – either detection or production; psychopathy predicts nothing.... The

Trust and Deception

good news is that these measures predict each other, which is kind of what you'd expect. Lie acceptability...does predict, so your ability to tell lies. If you think that the more likely you are to think lying is acceptable, the more practised you are at lying, which does seem to work. The big one from the literature, ability to deceive yourself, doesn't predict lie-ability. Doesn't predict this d' measure. It predicts your bias, so how credible you are. Hilariously, from our point of view, it predicts it negatively. So the more that you self-deceive, the less credible you appear. So against all of the theories in the literature.

It could be the case that we don't see correlations because we use bad measures, but in each case we use the gold-standard one and we find the expected pattern of inter-correlation on these measures. So we are picking up something that seems to be reliable and valid.

And I think one thing that's really important that we kind of forget to separate out when we think about lying is ability versus propensity. So I think it is the case that psychopaths tend to lie more, so they have a greater propensity to lie. That doesn't mean they've got a greater ability to lie. I think often we conflate those two things.

When it comes to the type of stuff that's going on in your head, the cognitive mechanisms involved in telling lies, we've actually got very little evidence for any specific cognitive ability. There's some evidence for executive functions, particularly your ability to inhibit pre-[0:51:00] responses – stop yourself doing automatic acts – and this is thought to relate to truth inhibition. If someone asks you a question, it's thought that the truth is the default response. I think that's probably wrong in a lot of cases, but anyway.... And so there's a very small correlation between tests of this executive function and truth inhibition. We wondered about your ability to control representations of yourself and other people. So we know, for instance, that people have a false consensus. There's a false consensus effect, so if ask you what you

Trust and Deception

think about abortion and ask you to estimate what the population as a whole thinks about abortion, your – you think your view is much more representative than it actually is, because we tend to hang around with people of like mind. And in general the idea is that we kind of use ourselves to model other people's views. We think that they're going to be the same as us. And so we had an idea that, actually, one of the reasons why we're so bad at detecting lies might be because we're able to step outside our own shoes. If we view that abortion is absolutely terrible, then we just don't believe that anybody else can actually be in favour of abortion, for instance.

So we used a big, mediated lie technique with a fair few people – that was about 60, and we did find that you're better able to detect lies when the other person has a consistent opinion – an opinion that's consistent with yours than when inconsistent, but as you can see, this effect is absolutely tiny. It's 2%, and the mean? 54%, which is wonderful. So, yes, it makes a difference whether the other person's opinion is consistent with yours, but it makes hardly any difference in the real world.

We thought this was interesting because we've got lots of ideas about how this self/other process enables you to kind of separate your views from other people's views – works in the brain. We've looked at this self/other thing in relation to perspective-taking, we've looked at it in relation to empathy for pain – stopping yourself imitating other people, and we use this technique for transcranial direct-current stimulation (sorry, that's my 15 minutes up; I'm going to be very quick) of the temporoparietal junction and it seems that we can boost people's abilities to step outside kind of their own mental states and represent those other people. We've shown that we can make people better at [0:53:31] [in]consistent trials and we use this kind of technique. Again, we're not making anybody any good. This is like 59%. So it's kind of interesting on a psychological level; useless on an applied level.

Trust and Deception

So, in conclusion, you can't detect deception. There's does seem to be some deception general ability. It's very hard to predict who's going to be a good liar. We don't really understand the cognitive mechanisms. I don't think we're going to get anywhere near to developing a good lie-detection technique. It's never going to work.

What I think's interesting are these credibility judgments. So we make these judgments in 40 milliseconds. They're incredibly consistent across people. And, actually, I think this is the main thing that matters. So some people walk into a room and everybody believes them, no matter what they say; somebody else can walk into a room and everybody thinks they're lying before they've even opened their mouth. We can use these techniques to measure this reliably, and I'll stop there, 'cause my time's definitely up. Thank you.

S2 Great. Thank you.

[Applause]

Fascinating stuff.

END OF RECORDING