2._Collective_Intelligence_and_the_Roots_of_Social_Cognition-Mobile

Time to start again. Welcome to the roundtable on collective intelligence and social cognition. Why social cognition? Why the social mind? Well, a certain way to think about collective intelligence, claims about it, usually claims about the wisdom of crowds or groups being smart. And some people go so far as to say that on some occasions groups are smarter than their individual contributions. Well, claims about collective intelligence usually are part of broader discussions about social cognition and the social mind. In other words, the ability of individual people to interact and do things together.

Now, it is a pleasure, indeed an honour, to welcome two among the most influential leaders in the society of social cognition. Here with me we have Chris Frith and Robin Dunbar. Now Chris is professor emeritus of neuroscience at the Wellcome Trust Centre for Neuroimaging at UCL in London, a regular visiting professor in the Interacting Minds Centre at the University of Aarhus in Denmark, and last but not least, last year Chris joined the Institute of Philosophy in the School of Advanced Study at the University of London, as an honorary research fellow. Now Chris has done wonderful work on the applications of functional brain imaging, the study of social cognition. In 2014, he and Uta Frith were awarded the Jean Nicod Prize in Paris.

Robin Dunbar is professor of experimental psychology at Oxford University and the leader of the Social and Evolutionary Neuroscience Research Group. A world-renowned anthropologist and evolutionary psychologist, he's done work on the relationship between evolution, cognition, and behaviour, in particular.

Now, Robin is starting with a presentation that will have like a 15-20 minute presentation. Then each of them will have a chance to reply to the other with a quick round of comments and then of course we'll open the floor for questions.

What I really want to, really-- I don't want to present a negative view. All I'm doing really is saying this is one of the challenges that we have to overcome in order to create a large-scale collective intelligence, or even a large-scale democracy, of the kind which the previous two talks about.

So you might think that your social world looks like the picture on the left there and it's potentially enormous and huge. Indeed, that was the promise on the tin can made by Facebook. Indeed it was made by email 20 years before, I think, that these new digital technologies are going to open up vast new social vistas for you because you're going to have friendships all over the world.

The reality is, your social world doesn't look anything like that. It looks much more like the one on the right. That's to say, it's very small and it's very clumpy. Question is, why? The answer goes back to what's known as the social brain hypothesis, which is an explanation that was produced in the late '80s for why primates should have unusually large brains. And the explanation offered was that they live in very complex social systems and therefore they

basically need a very big computer to handle that constant dynamic complexity of relationships in those social systems.

And that's a consequence of the kind of bonded social systems that primates, in particular, have. They exhibit these kind of intensely bonded relationships universally, across the whole order, at least for the Anthropoid primates, and are very different in that sense to the kind of social sociality found in other groups of mammals, with a few exceptions. Some of the exceptions, a group like the horses, the horse family, the dolphin family, but we worry about those.

For primates, the key issue is that if you plot social group size, as I've done here, against a measure of effectively the size of the brain, and particularly the neocortex, then you get what turns out to be a series of greyed lines, in which species lie quite close to what seemed to be quite strict grades in which if you live in a big group you have to have a big brain. But the grades reflect as you move-- I've shown two monkey grades there in blue and the ape grade to the right-- as you move from left to right, you're going through increasing social and cognitive complexity as the species involved support more complex, socially complex groups.

Well if you plug humans into this, you get to the ape line. That is, you get an estimated group size for humans of 150, which is now known thanks to Facebook as a well-known emporium of science as Dunbar's number. Question is, do we actually live in such small groups, given what we actually have behind us after the window as it were. The answer seems to be yes.

So on the left is just some casual observations of organisations, which occur in that sort of size as the typical size. And one to 200 is really the outer limits on the prediction round, that prediction of 150 for humans, and actually reflects the size of typical ranges variation the size of human social networks. On the right hand side, various quantitative studies, which show the real data behind this. All of these, some of them and looking at top down at the way the world is structured as it were, others are looking bottom up at your personal social network. All of them pinpoint this figure of 150 as the typical value. Again, with lots of variation, individual variation, around that. But still, this figure keeps coming up.

So it looks like there's a natural size, if you like, an actual structure to human organisations. Part of the reason for that comes from the fact that, aside from the cognitive constraint, which is what I've sort of really been talking about up till now, the brain size effect, it is a timing constraint. That we simply cannot get round and talk to people. And these are data, which Dezecache collected in Oxford some years ago now, looking at this natural size of conversation groups. And basically what's plotted there in the dark upper lines, the relevant one for the moment is the number of people engaged in conversation together, plotted against the size of the social group they were actually in at the time.

You can see conversation group, the number of people paying attention to the speaker, including the speaker, rises to an upper limit of somewhere around about three and a half to four people. You simply cannot get more than four people in the conversation. If you have more than-- when the fifth person joins, I will bet you a case of champagne that within half a minute-- sorry, Mattia will bet you a case of champagne-- within half a minute that will be

two conversations or it will become a lecture. The only way you get a bigger conversation is to have formal rules of lecture if it's with a chairman and everything. If it's completely unstructured and natural, it will become two conversations very, very quickly.

And here is sort of a reflection for that. This is actually looking at the rates of contact that people have with the members of their social network. It's the rate per day for each member of the layers of their social network, which I'll show you in a minute. Something like 40% of your total social time is devoted to just five people in your social network. It turns out it doesn't matter whether you measure this as actual time invested or whether you measure it as the emotional investment that you have in that relationship. The two correlate very tightly indeed.

About another 20% is devoted to the next 10 people. So 60% of your total social capital, in effect, in your everyday social life is actually invested in the 15 people who are closest to in your social network. And this actually creates-- you have a kind of social signature, which looks like a kind of fingerprint for yourself. The way you choose to distribute your time within these averages is very characteristic of you and quite independent, stable over time and quite independent, of who's in the network. In other words, if you throw some people out and bring some new people in, you will slot them into the same time slot as the person you've just thrown out.

So this goes back to the fact that of course this is all being done face-to-face. Your social network in real life looks something like this. It consists of a series of layers. The innermost layer of five is this one that you invest most heavily in. Then you've these layers that go out and they have a very distinct scaling ratio, so each layer is three times the size of the layer immediately inside it. Goes out to this figure of 150, what the red line is. That's a real boundary in terms of your relationships with the people involved.

We know there's at least two more layers outside that, one at 500, one at 1,500. But the distinction between the 150 and the 500 is the 150 are friends and family, the 500 are acquaintances. And that distinction affects dramatically how altruistically you will be prepared to act towards any member of that group.

Turns out you might ask, in the light of the previous talks, of course, whether online digital media have changed all this. The answer is, not one iota. The average number of friends on Facebook is around 150 to 200. We've just done up two 2,000 person structured random samples of the UK of 18 to 65 year olds, asking them to tell us what they have numbers of Facebook friends. Both samples produce means of about 150, obviously with a range of variation around that.

It seems that whatever is happening on Facebook doesn't really change how you interact with the people in it. Yes, you can sign up large numbers of people if you want, and a very small number of people do, a very small number. Something like only 5% of people have more than 400 friends on Facebook.

That's large enough when Facebook has a umpti-billion members for everybody to know somebody and that creates the urban myth of everybody has 1,000 friends on Facebook. But

the short answer is, they don't. And if you look at the traffic that goes on, who's talking to who, then what you get is, these layers reappear very nicely. Here they are for Facebook, in fact, we have three different Facebook data sets. This is just one of them.

We also looked at Twitterati. That's to say, the followers on Twitter talking to each other on a Twitter account. And you get exactly the same numbers. They are so close to the offline layers, it's just uncanny. What's also interesting, as it turned out, and we'd always joked about this before, there is actually an inner layer exactly where it should be at 1.5.

As an aside I might just say it's 1.5, because this is your best, best, best friend. This is because blokes can only do one at a time. They either have a girlfriend or a best mate, bloke. Girls can do two. They have a boyfriend and a best girlfriend, what's known in the trade now as a BFF, a best friend forever. And we didn't coin that.

So it looks like the online world is simply mimicking what's going on in the real world. Just as an aside on the previous two talks actually, one of the things I was told by some of the people who are working on the Arab Spring uprising, particularly the Egyptian data, they said yes, yes, Twitter was very important in orchestrating it. But it worked because of a small group of people who had face-to-face relationships with each other. They were the ones that really put it together and made it work. All Twitter did was kind of coordinate the masses, if you like. But the real work that made it happen was done by a group of people who've known each other for a very long time.

Just as a reflection of the limits of information flow that you get, and this is a model output of the likelihood, essentially, that you would find out some piece of information that was seeded somewhere in the network. And it's simply looking at the consequences of a piece of information being seeded in the outer reaches that you might hear about it on the basis of the frequencies of contact with different members of your network that we found in our data. It's the data I showed you earlier.

And what seems to happen here is really once you get above about 150, the likelihood of hearing about something, it simply stabilises. It's something seeded right out in the outer reaches of the 500, or the 1,500, or the 5,000, the next network out, of the next layer of the network out. Your chances of hearing about it are very, very small by face-to-face. It's not too surprising. What's interesting, if we use the same digital data contact rights, actually, it does continue to increase. So one thing the digital world does do is allow information to flow much faster because you see it. But it's because it's overriding or bypassing the face-to-face constraints.

I might say just two interesting asides here. One is that Plato-- well, Aristotle and Plato got these numbers right, which is just evidence, I think, that in two and a half thousand years since the Greek philosophers, science hasn't actually improved. Plato-- Aristotle said you can count the number of real friends you have on the fingers of one hand. That's the inner layer. And Plato said the ideal democracy size is 5,000. That's the next layer out after the 1,500. But what's interesting is, these are also the structure of modern armies.

Modern armies have evolved these structures over time as a matter of trial and error. And they keep the ratio going, beyond 1,500, 5,000, 15,000, 50,000. I think 50,000 is the division. And no matter what names they give to it, all modern armies have this shape. It is all about control and management of troops on the battlefield. The only layer they don't have is the five layer. But that's the army of the future.

The new face of warfare, which is no longer lots of people standing on battlefields, firing at each other. Turns out that five is the standard size of terrorist network, terrorist cells, and five is the operational unit for special forces. And both of them have those numbers for a very good reason. That these guys have to be deeply, deeply bonded to each other so that they stand by each other under all circumstances. So it's kind of interesting. Anyway.

The consequences of this kind of natural structuring, and I say again, it's a consequence of two things running together. One is this cognitive constraint. We kind of know it's a cognitive constraint. We now have loads, well, about six or seven imaging studies have been done now, all of which have very nice correlations between individual differences in the size of social network measured in various ways and the size of particular bits of the brain. And then the other side, this time constraint problem.

But one of the consequences of this, I think, is self-obvious in that most modern management structures are basically a complete disaster. They try to impose a hierarchical management system in a context where it can work because that's how the military get it to work. But the difference is, the military imposed draconian discipline on the system. That's the only way they can get this enormous hierarchical structure to actually work and function reasonably efficiently.

If you look at most disasters, and here's the latest only with [INAUDIBLE] and share price collapsing at the end. It's all a consequence of fractionation, almost certainly, within the system. That once the system gets quite large, it's very easy for sectors within that, some people are now referring to them as silos, to simply hive off and do their own thing against the interest of the rest of the country, the rest of the group. It's also very easy for the Nick Leesons of this world to sort of hide what they're doing from the rest of the system and end up bringing banks down, as happened in that case.

And here's part of the problem. It's what's been referred to as the water cooler problem. It was actually, I think, first identified by Wil Gore who set up Gore-Tex, probably the most successful medium-sized business in the world. And what he basically said, having worked for DuPont, the big American chemical giant, previously, is that organisations become dysfunctional very quickly once people don't know the members of the group that the organisation they're working with.

And he pinpointed a figure of about 150. I only wish he'd consulted me. Unfortunately, did this about 30 years ago. But I tell you if he consulted me on this, I'd be on a very big yacht in the Caribbean with Chris Frith dripping grapes into my mouth on one side, and Mattia here dripping pina coladas on the other side, from the consultancy fees. Because Gore-Tex is widely thought to be, among business organisation theory people, one of the most successful medium-sized businesses in the world. It's been incredibly successful.

But what they identified was this problem that if you have too many people and they don't actually know each other, a we versus them culture very, very quickly builds up. And it's simply basically because people can't get around and talk to each other. So what Gore did when he set Gore-Tex up, literally in his backyard, was to insist that you had completely separate, very small factory units, and sometimes these factory units are literally on the parking lot next door.

So as the company increased in size, as business boomed and they want to increase, they didn't do what all the multinationals did, which is just make the fact is bigger and bigger. What they did is insist that no factory ever be bigger than 200 individuals, but preferably 150. But they're completely self-contained. It's what they refer to as their flat lattice management structure because each of these units is completely autonomous. It's just the board, the family board sitting above it, gives them the objectives and goals, as it were, one by one. The factory gets on with it.

But because the unit is so small, you don't need any management structure, formal management structure, so that everybody has the same label on their jacket. It just says Gore-Tex associate. Because everybody knows who the factory manager is, who the accountant is, who makes the sandwiches, et cetera, et cetera.

And secondly, you have this intense loyalty to the community, that little community, which makes people then work together and collaborate and be more altruistic in negotiating difficult problems that they have to overcome. So I strongly recommend this management structure to all of you. I think it's absolutely brilliant. Nobody does it and I just leave you with a reminder, three areas which I think potentially suffer from this enormously. One is the NHS, which is insisting on building ever bigger units. The other is Whitehall, for goodness' sakes. It's a disaster area.

An the third I just point out to you is schools. Secondary schools are already too big. The average size of secondary schools, state secondary schools, is about 1,500. It's miles too big, miles too big, even for 18 year olds. And there are plans afoot now to build about five mega schools with 3,500 kids in them. It's going to be a disaster area. Best of luck.

Thank you. Robin, thanks so much for a wonderful and very thought-provoking talk. Now Chris on the, on the question, how can the group be more intelligent than the individual.

I'm not going to use any slides, at least as far as you're concerned, because I made an individual decision that this was a bad idea. But I am using slides that I can see but you can't. And I'm going to talk about, is it the case, and when is it the case, that groups can make better decisions than the best individual within that group.

And there are various sayings, like we know about the wisdom of crowds, which suggests this is true. And there's also the folly of crowds that suggests this is not true. And in fact, Bahador will be telling you much more about this this afternoon, so this is a sort of warm up for him.

But I think, as was implied in Robin's talk, in human society we actually very rarely make decisions on our own. This is true both at the family level is two and a half or five people right up to the national level. But the question I want to ask is connector decision-making actually a good thing? Is it wise or is it folly? And if it is a good thing, a particular challenge for the study of collective intelligence is to find out when group decisions are better than individual decisions and when they're worse.

And there are two basic, and I'm now talking a very basic level of decision-making, there are two fundamental reasons why a group should be able to make a better decision than the best individual within the group. And the first concerns the stasis of the information that each individual holds. So they may have some knowledge which will help to make the decision, but it may be very imprecise. So an individual on their own might be able to make the decision, the good decision, slightly over half the time.

But if all the people in the group are functioning at this level, so they're all slightly better at chance in making the decision, because information is imprecise and noisy, then by combining their signals, as long as they're all different in the nature of the noise in the imposition, if you combine the signals together the noise can be reduced and the information could be made more precise. So that's one reason why a group should be able to make a better decision, even the best individual. And this, of course, is a fundamental property principal in statistics it's the basis of signal averaging and all sorts of other practical techniques.

And, in fact, this kind of information integration is widely seen, for example, in fish. Fish are very good at making group decisions. The example here is if you have a very weak signal in the water, that is a smell or trace of some chemical, which points in the direction of food, each individual fish may not-- the signal may not be precise enough for them to move correctly in the right direction. So they're all slightly moving in that direction, but there are milling about randomly.

The fact that they're in a school, where every fish's behaviour is determined by its neighbours, means, actually, the noise is reduced and the signal is increased. And the school of the fish can follow the trace, but the individual fish cannot. So that's an example of how this collective decision-making might work.

There's a second principle, which is slightly more obvious. If the individuals in the group all held different information from each other, then a more informed decision can be made. There's more information available to the group then there is to any one individual.

And this, again, is seen widely in nature. This is how honeybees choose a new nest site. So when they have swarmed, they're hanging in this tree somewhere, and roughly 100 scout bees will get out in different directions to look for new nest sites. And they find various sites and they assess them for how good thing are, which is things like size, not being damp, having to hole the right distance up the tree, and so on.

So they have an estimate of how good the site is and then they return to the swarm and do their dance. And the dance indicates, the orientation of the dance indicates, where the site is. And the length of the dance represents how good the site is. And so you've got hundreds of

bees all coming back and giving different advice. But there's a sort of positive feedback mechanism so the dances that are longest are more likely to attract other bees to go and look at that site. And if it is indeed good, they will come back and give long dances, whereas if it's not so good, they'll come back and give short dances.

So over several hours, you might start off with five or six different site directions being indicated. But by this process of positive feedback, over several hours, it eventually is reduced further and further until the vast majority of the scouts are all pointing to one site. And then, by some magical process, the hive, the swarm decides this is the right place to go. And they will go off to this new site, which is typically, is indeed the best site in the area.

So these are two examples, and they're sort of fundamental statistics about how group decisions could be advantageous. But of course, there's also, we know from statistics, that there are special requirements that have to be met for these group decisions to be advantageous. So first of all, the individuals making the decision must be independent of each other. If the noise-- if they're not independent, then the noise is not reduced because they're all having the same noise and it may actually be amplified.

And this, for example, may be the sort of thing that causes what are called information cascades, or stock market bubbles. So as I said before, each person in this making the decision about what stock to buy might have a slight chance of making the right decision. They know a little bit about the stock. And if their suggestions were combined independently, then they would probably make a good decision.

But if they announce their decisions public and in sequence, they can go off in a bad direction. Because the first person happens to make the wrong decision, the next person is influenced by the fact that the first person has made this decision and overrides what he thinks. And it builds up, so everybody goes off in the wrong direction. And I say that's called an information cascade.

Another reason why this may go on is when you see the other people making the decisions and you don't know very much about it, you will think they have inside information. That is why I should follow them. So these are examples of where the group decision goes wrong because the people are not making-- the information they're providing is not independent.

And I think this is quite interesting, in relation to what we were hearing earlier about what is the impact of the internet on this sort of thing. The internet, of course, is good in the sense that you can have many more people combining to make the decision, which should make it better. But it's also bad because the internet also causes the people who are making the decision to be much less independent because they're all reading what Twitter says or whatever it may be.

So these sorts of things, in a sense, counteract each other. So it's not clear whether internet-type decisions with thousands of people will actually be any better than low small-scale decisions with five or six people. And I would have thought also the internet will be more prone to produce the information cascades couples something one has to be very careful about.

Now the second thing we have to deal with when we're combining information to make across people to make a decision is the competence of each individual. We should obviously put more weight on the contributions of the more confident individuals and ignore in extreme cases the incompetent individuals. So there the problem is, how do we know who are more or less competent? It's particularly true in the internet where we don't know who we're interacting with.

So one sign, and I'm again going back to the nature in general, one side is the speed with which the decision is made. If you make your decision faster, it probably means that you know something about it and therefore perhaps your advice should be taken. Now this happens again in, for example, in homing pigeons.

So homing pigeons fly in small flocks, or even pairs. I should say that most of our work is done on pairs. A group means two. And interestingly, the homing pigeons pairs will navigate more efficiently than the best individual within that pair so there's an example of a group advantage.

And on each choice point one of the pigeons will make the choice of direction faster than the other and obviously, almost by definition, becomes the leader at that point. And it's typically the more experienced bird that becomes the leader because he's making or she is making the decisions faster.

So in this case, the fast decision is an indicator of competence. And this is often referred to as confidence. So I have to be careful. I have to make-- you have to remember there's one thing called competence and another thing called confidence, and that your confidence, hopefully, is a marker of your competence. And indeed, my friend James Kilner in Queen Square did a nice experiment where he showed that if people indicated their decision by making a motor movement, then the faster they made the movement, the more confident they reported themselves to be in the likelihood that their decision was correct.

So there's a good correlation between speed of making the decision and the confidence. And also observers could extract themselves, though you ask people who can look at videos of the movements. They would indeed correctly identify the more confident cases on the basis of the speed with which the actor moved. So this is an example of a marker of how confident you are. And you could then say OK, when we're working together, we should always choose the decision of the more confident person.

And this does indeed seem to work in these very simple tasks that people like Bahador study, you can-- which is trying to detect a weak signal two people doing it together if you simply get them to rate their confidence on each trial and if you choose externally the response of the more confident person, you will get this pair advantage. So this is an example of how confidence as a marker of competence can be used to help making group decisions.

And you can imagine, I imagine to myself how this might work. So if you're doing these very boring tasks where you have to detect these weak signals, occasionally you will not be attending and you will miss it. But and that will be true for both of you who are doing the task together but it's very unlikely that you're the lapses of attention will occur in both people

simultaneously, so as long as you can tell the other person when you had your lapse of attention you can actually do better than the best person on their own.

But there's a big snag to this situation because it only works if the two people are roughly equally competent. Obviously, if I'm more competent than you, I should put more weight on my decisions. And the question is, will this be reflected in the confidence rating. So these are the problems with the use of confidence as a marker of competence.

We all know, or I know, that I'm above average in competence and you probably all think the same thing. And apparently in one study 90% of university lecturers rate themselves as above average in as electros. So this is an immediate problem of the use of confidence as a marker.

The other problem is that we all tend to think roughly speaking when we're working together that that we're equal in our competence. So if two people of different competence work together, the joint decisions could suffer because too much weight is put on the choice of the incompetent person. You will say more about that.

So there has to be some way of calibrating your confidence and does that occur. And of course, in large groups you can get the problem that a vociferously confident minority can hijack the group, leading to decisions it would've been better if the whole group had somehow been tended to. And that might be a problem with the sorts of examples we were hearing this morning in particularly decisions made with very large groups on the internet.

There's a strange, as an aside, there's a study that came out I think last year, which is somewhat surprising, that said in this situation where you have a vociferous minority, overconfident minority hijacking the group, this can be overcome by having numbers of uninformed people who will somehow combat the vociferous minority. So that's a useful trick that one should bear in mind.

Now of course unlike other animals, I guess, humans, we humans, can manipulate the way we express our confidence so if I'm working with an overconfident person I might be able to overcome this by pretending to be more confident than I really am. I mean the reason I'm standing at this lectern is to make me appear to be more confident than I really am and it may even make me feel more confident.

And furthermore, unlike other animals-- actually it's controversial. So humans are very good at telling people how confident they are, how likely they think it is that they made the right decision. It is extremely controversial as to whether any other animal can do that. There is some slight evidence with rats and monkeys, but it remains controversial. Maybe they're doing it another way. So something I think is particularly unique about humans is we can actually introspect about how our mind is working and we can talk to other people about it. And this is necessary, both for expressing your confidence and for manipulating your confidence to use in situations where interacting with other people.

And I could give a whole lecture about this but this ability to introspect on how our minds work and discuss this with others is the basis of things like cumulative culture. It's the basis of the idea that we are responsible for our actions, where there's something called free will.

We all know about this and we talk about this and we in very early in childhood children learn to distinguish between things that are done accidentally and things that are done deliberately. I think this is the sort of thing that you can learn about through this ability to introspect and discuss what you're doing.

I have said nothing whatever about the brain, although that's supposed to be my speciality. So I'll very briefly say that there are now studies looking at the brain, brain regions I should say, that's very phonological, the brain regions that seem to be necessary for talking about how confident you are. And not surprisingly, this is in both the frontal cortex.

So the patients with lesions to the frontal cortex will be just as good at making a decision but they cease to be able to accurately decide whether it was a good decision or not. I say they're just, they're just as likely to make a good decision, but they no longer are able to introspect on whether they have made a good decision.

And in particular, it's the very, very front part of the brain known in the trade as problem area 10, which, and Robin will correct me if I'm wrong, it's believed to be the bit of the frontal cortex that is most expanded in humans compared to other primates.

Yes, and it's unique to Anthropoid primates.

Yeah. So that's of course I say that's just the phonology. What we really want to know is how the brain does it, which I'm sure Bahador will find out eventually. But the final point of course is what we can do uniquely as humans is we can have discussions like this where I am telling you about what we're learning about group intelligence when good decisions are made and when bad decisions are made. Then in principle, we should be able to use this knowledge actually to be able to actually make good group decisions. Well I have it yet

I'm sure you have very many questions on that certain claims, especially if they're both Robin and Chris made on the relationship between the digital world and forms of collective intelligence and collective behaviour. But before we open the floor for questions, I'd like to ask each of you if you want to reply quickly, starting with Robin, to the comments on collective intelligence to decision making.

Yes, I'm-- I agree. Obviously, I agree with Chris. And the, kind of, these cognitive processes and a lot of what we do simply wouldn't be possible. But I guess the question I keep inevitably coming back to, I mean we're kind of constrained always running experiments almost everything but particularly humans is actually doing it on anything but two that's at a time very often and I do wonder where the limits are.

I mean, it's perfectly logical, I think, that people in groups will always do better but that's the most is going to apply on the scale of maybe 50, maybe 100 at the most, but it's based on groups of much more sometimes I wonder how much that can be scaled up or whether we just hit buffers in terms of our ability to really think about what other people are doing because it all comes back to doesn't it sort of your ability to intuit what somebody else is thinking and how reliable their, what they're saying with that how actually, actually is and

how many minds can you actually do that with at any one time or in an interacting situation? So that would be my kind of picture things

Yes, I would guess I mean yes the problem is running experiments is that it becomes very expensive so typically you run in the olden days we only had one subject in our experiments and then if you have two subjects, that doubles the cost. And then you can see it becomes astronomical. There may be ways of doing it on the internet, which is becoming more and more popular with psychologists using Mechanical Turk or other platforms.

But in regard to Robin's question, I wonder isn't this perhaps why this method has developed where you have a representative so a group chooses a representative which is then sent out to make the decision with the representatives of other groups so you're constantly trying to keep as a relatively small the number of people who actually make the decision. So I guess that's a sort of hand-waving explanation of why this happens.

I would guess, just like with your conversations, if you're actually looking at a number of people making decisions together above about five, it's not going to make much difference. But remind me Bahador did somebody do an experiment on predicting what the numbers would be?

But what the problem is we got it The empirical difficulty of studying two versus larger numbers is that as the numbers change the our heuristics for how to collect or integrate decisions among ourselves change. Like one thing we have seen in the laboratories is that the moment we have three people write it down too, people to stop communicating the content of their information and they just count both.

So immediately they seem to think, OK there's an easier way to do this, so let's just get rid of the more difficult but perhaps more useful ways of doing this. So it's not very easy to adopt the model or to expand the model for two people to, let's say, three or five or 20.

I had a question for Robin because I've always been fascinated by your idea, which you didn't particularly mention, is the importance of gossip. And the gossip I believe is actually very-- there is something I didn't talk about either, which relates to this confidence story, is the idea of reputation. So that some people, but this particularly applies to the economic trust, so people will build up, some people have good reputations. That is to say if you invest your money in them you get it back. And other people have bad reputations because if you invest your money in them they steal it.

And gossip, and it applies to any sort of contract you have with people, like your plumber or whatever it may be, gossip is extremely important for transmitting, I think, the reputations that people have. So in a small village you know who the trustworthy people and your don't, and you know who the untrustworthy people. But what happens, this seems to be a very important part of social intelligence, what happens with big groups and on the internet?

They lie. The answer, the short answer is yes because gossip, or indeed the verbal transmission of information by language, is actually crucial in allowing, or was crucial, in allowing humans to up the size of their groups from the roundabout 50 as an average species

group size, which is the biggest, most social species of primates, monkeys and apes, ever have, up to this kind of 150 figure for humans. And in lots of different ways.

I mean there's no doubt language has been very important in making that possible and reputation is clearly important. But even monkeys and apes and probably a lot of other species rely on reputation

Fish, yes.

Fish. But the difference, as you rightly point out, is in their case it's the reputation you see. Whereas for us it's the reputation you can promote for yourself. And that automatically brings the downside to it in that you can fill the ether with black propaganda about how wonderful you are and how absolutely atrocious everybody else is.

So, yes, it works. But I think it then comes back to the fact you will also probably have to have very good non-verbal cues, including familiarity. And a lot of what these interaction rates we see between dyads in friendship networks is about building up that confidence just by being with them and seeing how they behave. So it's not a completely foolproof problem. But then nothing in biology is. There's always a downside to everything that you invent.

On this note, I'd like to open the floor for questions. All right, let's start with a question.

Hi. Thank you. Yeah, I, hearing that and we do better in groups of sort of three or four and that vociferous people can hijack groups, all that sort of thing, is incredibly familiar to me. I've spent my life doing qualitative research and public dialogue trying to, kind of, work with those things group effects. I guess my question really is about if competence is so hard to achieve in sort of heterogeneous groups where people don't know each other, where there's no trust and reputation, where we defer to each other's decision-making, even if it's bad. How do you have any kind of way of measuring whether decisions made in groups are in any way competent in any way that can then be relied upon in a democratic process? Perhaps to return to the earlier session where we were talking about whether things are legitimate, whether intense enough, whether they're deep enough, whether they're wide ranging enough. How do we smush together all these complicated things to come up with anything we can trust when it comes to the public decision making?

There were two points that you were making there, one of which is, which is perfectly correct, that we devise situations in the lab where we know which is the right decision and which is the wrong decision. Whereas in real life, you often don't, or at least not for a very long time. So this makes such studies extremely difficult. And I guess it's why this, it relates I guess to this idea of accountability that people who made the decisions should somehow be visible and then they are accountable.

There's one story that the whole structure of the British Civil Service is designed so it's impossible to work out who actually made the decision. And this is presumably why. I mean there are various rules that Bahador might talk about, about how you might conceivably help to make committee decisions better than they often are. But what it's not just a matter of knowing who's competent and who's competent at what. There's also a matter of everybody

actually having the same goals, which is sometimes not the case, or often not the case perhaps. And I'm not sure that I have any sensible things to say about it. But these are the problems that we should address.

The thing I might just very quickly just make the comment that in real life a lot of the decisions we make don't have a right or wrong. Politically, at the level of government, they don't have a right or a wrong way solution, as it were. They're just things we have to agree about, like all driving on the right-hand side of the road or the left-hand side of road.

And in small-scale societies, which have this typical community size of about 150, that is done as, essentially, as a democratic decision. And our problem, in some sense, is once you're in the scale that we're in, and remember our parliamentary democracy here was sort of created around the time of, well, before Magna Carta probably actually, when the population was very, very small indeed and we're still using the same structures but they've just expanded out of all proportion, probably to the point where they effectively dysfunctional.

Our problem, if you like, I think for the future is how we bridge that gap. We know what the constraint is, in some sense. So how do we work round that? Again, in biology, as in evolution, as in the internet, all things are possible. It's just actually understanding what you need to-- what the constraints are on how to bypass them that is, in the end, is the problem.

Can we invent new ways to go beyond the limits that exist and are hard-wired in the brains, and yet we can maybe produce things, like academia institutions or that are somewhat outdated in many dimensions. But can we renew this knowledge institutions and can we invent new sort of knowledge and acknowledgement institutions that can help identify the best possible solutions and the best, the most promising conversation so that we can keep the ball rolling and look for new innovations? , Because, you know, bacteria for instance are very good at collectively including different species, sharing information on the best way to fight antibiotics. And so they have coordinated efforts on skills that are just unsinkable compared to the sort of discussion we're just having. And I think that there is ways of building alternatives that we have to explore and maybe that's what we're here for.

Thanks.

Yes, the way ants and bees make decisions and as you say the bacteria as well, is very interesting. But I think I really want can see why humans are different because as far as I understand the ants and bees do not have a confidence or competence way to applied to the different individuals it's always equally and of course, they have typically

[INAUDIBLE]

But I'm not sure whether that reflects competence. Yes, I think we already have a more sophisticated way of doing it. There's a very interesting paper that came out last year or so where they have worked out the mechanism by which bees do their interaction, make their decision in great detail, so there's not just the positive feedback. There's also what equivalent of lateral inhibition. And their sort of last sentences that the mechanism which bees use to

make group decisions seems to be identical to the mechanism which the neurons in the mammalian brain interact to make decisions.

So the bees working together can simulate the mammalian brain. So the obvious question then is, if, so in that case, mammalian brains working together, and that includes the humans, should really be able to do something exceedingly exciting. But as you say, we don't quite know what it is.

Robin.

I'll just make two comments, actually. One is, yes. I mean, I agree with you. What some of these lower organisms are capable of doing is amazing. On the other hand, I think my impression is, with my zoological hat on anyway, is that most of the decisions made by bacteria, or by species like ants or bees or the like that have very large communities, is quite rigid in what they do. And the bottom line in the end is, and this may be the trade-off we have, the reason we have big brains is precisely to be very flexible. That's what big brains give you in evolutionary terms.

So my observation might be simply that, well, yes, as yet ants and bacteria haven't built space stations or the works of Shakespeare or all these other things we do. And it's in those kind of cultural things that all these innovations that we've created historically have come to be. I mean, the worry is, there are some quite worrying statistics on this actually because the, the rate at which the world population has increased over the last 10,000 years seems to be driven by the frequency with which new technological discoveries are made. It seems to open a door, which allows the population to increase dramatically. But the rate of discovery is declining. So we're heading for a sort of collision course between increasing population growth and declining ability to create innovative technology to allow us to support bigger things, populations. And that's, you know, we do have to worry.

The other point I might make is, and it bears back actually to the question Chris was asking earlier-- one of the things they've done is model the efficiency of networks, of people in this case, agreeing, effectively agreeing, on some common cultural icon for us to solve a problem. And one of the things that surprised us that came out of that, and it will probably surprise you, may disturb you, is that these networks, communities, are much more efficient if elites arise that control the system. It increases the efficiency of communication and the flow of information around the system and everybody actually does better. Now at the moment we've only done those on a fairly small scale. I mean, our sense of small-scale community system, so we don't know if that scale is up or not.

Again, the problem is, and it seems to happen if in these systems adjust these are agents in an agent-based model doing their thing and just automatically they create these elites. You don't get, you get very highly structured networks. It's not democratic at all. What we don't know is whether that scales up to the big scale because you can see the obvious problem, and again I emphasise nothing in biology comes for free. There's always a downside to it sooner or later. So it's very easy once you have an elite for an elite to control the system for their own purposes for us in the collective good. And the problem may be that an elite works well when it's a small, a very small, scale community because they then have obligations to the members

of the community and the members kinship obligations them. But once you become unknowns, effectively strangers, you lose that control.

One of the nice examples of that is the Hutterites. So the Hutterites live a communal farming life in the Dakotas mainly in the USA. They've done, they moved from Czechoslovakia somewhere like that in the early part of the 19th century. They insist on splitting their communities, these communal farms. The farm is run as a communal venture right and their decisions are made actually democratically. They insist on splitting the farms once they get to 150 people, which is the smallest size they can really split them at and get two viable farms. They will hive off a group and send them off to create a new daughter farm and off they go.

And the reason they do that is they say you cannot run the system by peer pressure and democratic processes alone once it gets above about 150 to 200 people. In other words, when you say to somebody I really didn't like what you did Saturday night after the barn dance, if it's under 150 people they'll go, oh my goodness, we're terrible, I shan't do it again. If there's more than 200 people, they just go, who cares, that's your problem.

And that I've heard that comment so often in those sorts of contexts with at this hinge point of about 150 people. Suddenly you're into a world of strangers and you don't care anymore. And self-interest then becomes more important. Then how do you create mega communities? Because the same social cohesion as you're 150. And that's the question.

We have two more questions from this side of the room. Yeah.

May I, out of interest, query or probe Colin Blakemore's, I'm sorry, Robin Dunbar's statement that the standard size for a terrorist cell or specialist forces, a special forces unit, is five. My recollection, perhaps faulty was that in the UK Special Air Service, SAS, the standard module is four, or perhaps for a slightly larger or small group, eight, which of course really halves into two fours. You appear to be nodding to this. If you agree, would you think it has anything to do with the conversational cohesion breaking down above four? Because the SAS does seem to pride itself on close-knit cohesion in small group size.

I'm happy with four. All I'm saying is that that inner core, which on the general number pattern will be about five, that seems to be the target operational unit size. I think it actually has much less to do with my guess is and I'm we've not looked at that and I don't know if the special forces have looked at these things very carefully. They may well have done.

But my guess is it's not about what you use language for, in gossiping. It is just being there together constantly so that you just know how the other person thinks. You don't even need to ask how they think. You know what they're going to do then you know how they're going to respond. And this is true even of football clubs this point. pointed out by a number of people you can tell when a Premier League football team is getting on socially and nobody's stolen nobody else's wife or girlfriend at the weekend drinking session, because they play better. They just know, they spend all this time with each other, they just know what each other are going to do. They don't even have to look.

And these inner core relationships, the five and 15 layer, seem to be really built round that. Language is much less important for that. It's really just hanging out together doing stuff together and particularly doing physical exercises together, which trigger part of the pharmacological mechanism, which seems to be deeply underpinning the whole of private sociology. In particular, your bonding process and that's the, primarily, the endorphin system.

And doing things like marching or drills or singing or what have you, dancing Does the SAS dance? They should do. Is extremely good at triggering the endorphin system and particularly if you do it in synchrony. So it's just the-- it's providing the opportunity to constantly, constantly reinforcing this deep bonding process. And language is just laid on top of that. Language, I think, because much more important for the outer layers.

Chris. This one that my friends is have done a very interesting simulation of different groups playing various games and they found, at least in these terms, which is to do I suspect with specific interactions rather than just the bonding, that three is a very bad number because you tend to get two versus one and that four is much better than three and two is better than three but three you should avoid.

Time for a very short question. Last question, [? Orestis. ?]

Thank you. So I think that you try to hack into Condorcet's theorem, which says that if you have a pool of people and they're independent of each other, and equally weighted, and all have above chance percentage to make a right decision, the bigger the pool then the more likelihood you have you will get to the right choice. And you say, let's see how we can now hack into that by converting faster to the right decision by seeing who is more likely to provide the right decision and give them more weight. And you suggested that an indication of being likely to give a right answer is speed and reported confidence.

And I thought that was interesting because I always thought that it's, it could be how much information the individual has consumed. And obviously, this is something that we can now monitor with internet. However that alternative creates problems like, who decides what information is good and also problems of privacy although we shouldn't expect that bees have much privacy. But I wonder what you have to say.

I think I mean certainly the idea that people should be weighted on the basis of their competence is clearly sensible. And as you say the problem is how do you, how do you know how competent they are. And at the moment, I mean, either you test them, which usually you don't have any access to that, and confidence is a sort of proxy to this but it's a dangerous proxy. Interestingly, there are some-- but again people, the other members of the group will get some-- you can start to learn whether people are being overconfident or not.

And there is an interesting study of witnesses in trials where a confident witness will tend to be believed by the jury but if at some point he is shown to be wrong then it now goes against him. So you can very rapidly switch from accepting that someone is competent because of their confidence to decide to knowing they're overconfident and you shouldn't trust them. But there have to be mechanisms for enabling us to sort of make the switch, if necessary. But I have no idea, it'd be interesting, I mean there are interesting techniques developed on the

internet about reputation like you fill in things after you bought it on eBay or whatever. But I wonder if there will be similar internet tricks for working out how competent people are. But I have yet to see them.

Robin, one last reaction. I apologise with those who raised their hands. I'm sure Chris and Robin will be happy to answer questions over lunch. Now before we break for lunch, let us thank again Chris Frith and Robin Dunbar for [INAUDIBLE]. Thanks.