Seminar Notes On ‘Predicting The Future’.

Abstract: Job descriptions are poor indicators of useful knowledge that lies hidden among members of an organisation. Dr Bernardo Huberman shows us ways of identifying those who have knowledge relevant to particular problems, how to give them incentives to produce reliable answers and how those answers may be aggregated in a timely fashion. The mechanism includes a 'peer to peer' system that identifies experts, a novel privacy-preserving way of conducting surveys and a method for establishing sets of individuals with the same preferences without revealing what they are.

In the second session Dr Huberman presents a new economics-based technique for predicting the future of uncertain outcomes that uses small groups of people participating in an information market. The experiments carried out at the Hewlett Packard laboratories, Palo Alto, USA, show that this technique vastly outperforms both the 'imperfect market ' and the best of the participants.

(These notes constitute a synopsis of the presentation followed by discussion. Dr Huberman's slides are referred to in the text and a full set can be found at the end. Because of sound recording difficulties, only Dr Huberman is referred to by name).

Director of Complexity Research : Eve Mitleton- Kelly
London School of Economics
Houghton Street
London WC2A 2A

Presenter : Dr Bernardo A Huberman, HP Labs, Palo Alto, USA.

Compiled For The L.S.E.
by Geoffrey J.C. Higgs 31/1/03

Session 1

Professional people work within a discipline or community where an individual's working knowledge and ability to solve problems depends crucially on others. If we look at any professional community what we see is people working on perhaps different aspects or parts of problems that are common and communicating and publishing their conclusions so that progress can be made and recognition given. Some two years ago I and colleagues looked at co-operative problem solving, and set up a series of experiments that simulated this activity in a simple way. These involved a number of networked computers that collectively could solve some complicated problems which were essentially divided into different parts so that different bits of computer software in the system could work on different aspects and the results communicated between them. What was demonstrated was that when the right software programs were linked up and made to collaborate there was incredible speed-up in the problem solving. The fact that some of the solutions were erroneous and had to be discarded was an added bonus in that it mirrored what happens in real life.

This initial work led on to a study of what in the United States are called 'communities of practice'. Traditional organisational models are very simplistic in that they merely tell you who works with who and who reports to whom, but a close study of communication by analysis of e-mails and other forms of communication reveals networks that do not necessarily correspond to formal hierarchies, but constitute and informal community of interaction. What we set out to do was to characterise such communities in both a theoretical and practical way by looking at how they evolved and effectively solved particular problems that the formal hierarchy could not. Analysis of e-mail traffic inside large organisations can be used to validate some of the predictions about the evolution of such communities and we are now carrying out work in the Hewlett Packard research laboratories (HP), in which access to e-mail logs in the company enable us to see both how information propagates and which people are key in the propagation of ideas within the organisation. A basic intuition behind this work is that there is an intelligence that transcends what is in an individual's head that is distributed in some way within the whole community. The success of an organisation, whether it be academic, governmental or business driven depends on its ability to extract, analyse, aggregate and act on information quickly. The problem for a large organisation generally is to gain access to data which may be known only to people on the edges of that organisation. This situation can be seen in universities where students are often the first to pick up new knowledge or in a business organisation where the sales team will be the first to know what is going on in the market when the chief executives do not. A classic example of this was the near collapse of the CISCO company which was a darling of the Internet 'bubble'. Whilst the sales people were detecting that orders were beginning to slow down the C.E.O. was saying that he was going to hire every unemployed person in the area because he was totally convinced that the US economy was doing better than the media were predicting.

Data is, however, only useful if it is aggregated in interesting ways and people display a variety that is not easily modelled on a computer program. What are needed are new mechanisms which are designed to tap into this kind of social knowledge which take account of this variety. Individuals need information for themselves or for the company that is accessed quickly and is true. But there are a number of important issues concerning this kind of information gathering. One is the issue of privacy.

People do not like to reveal information that might be compromising and they are often ‘strategic’ in that they may not give direct answers for a number of reasons. In the United States and even more so in Europe there are stringent laws protecting peoples’ privacy and what are needed are mechanisms which allow the information to be revealed yet allow the individual to withhold identification.

A further important factor is that people are often ‘irrational’ in an economic sense in that given several options they will not always choose the one that is maximal for monetary reasons. People differ not only in their opinions but also in the ways that they are prepared to reveal information and cannot be treated as if they were just another asset in a company’s resources. Care has therefore to be taken in designing compensation mechanisms for example. In the second part of this presentation I will describe a methodology in which people are invited to reveal information in return for money. I strongly believe there is ‘cheap talk’ in which people will tell you anything they fancy, but if there is money riding on the outcome they become much more careful. If I ask you whether it is going to rain tomorrow you will give me an instant opinion, but if I say that if you are right I will give you $100 but if you’re wrong you will have to give me $50 then you will probably seek extra information to support your answer. That’s the assumption behind much of what I am going to tell you about in the second part of this presentation.

It is often the case that individuals have talent or knowledge that their job descriptions do not list and a great deal of money can be saved by eliciting such information. At the end of the year 2000 everyone was talking about the ‘millennium bug’ and lots of money was spent trying to predict how badly it was going to affect this or that company and how it might be fixed. The typical reaction was to bring in consultants in information technology and yet there might have been some person or persons within the organisation that could have assured the management that everything was going to be OK. One mechanism that we have designed, tested and deployed allows such people to be found.

In the past the problem with establishing a data base which carries informal as well as formal information about individuals is that firstly there is no direct incentive for people to provide a set of profiles about themselves and secondly that such data quickly becomes out of date. Moreover there is the problem of overkill. Suppose I wish to gain some informal information. For example, I am intending to travel to China and require advice about finding a particular kind of food. I cannot just send a memo to every individual in the organisation requesting information that might seem trivial to other people. On the other hand why should the person that has the information I require make it freely available. It isn’t something that the company needs to know and yet it might be important to me.

What was needed was a system that directed questions to those that were best able to provide answers and yet was one that preserved their privacy; a ‘peer to peer’ system with a different kind of architecture. The assumption behind it was that every person in the organisation had a personal computer on their desk in which a piece of software could be installed that continually made a profile of the user’s interests, not by asking them questions, but by data mining information that was communicated or accessed. As a person goes through the day, perhaps checking the weather in London or the price of tickets for a vacation so this piece of ‘client’ software builds a profile by selecting keywords. Now I don’t know what your reaction here in Britain would be about that procedure. In the States people get nervous but we have to remind them that the company owns the data by law, which it has to do to protect itself against

litigation. The data however stays inside the computer and though it could be extracted by the people that set up the infrastructure no other co-worker can access it without a password.

The formal and informal profiles evolve in time as the user's interests evolve. When someone wants, for example, an answer to the question, 'Has anyone dealt with Microsoft lately?' the person puts in a request which goes randomly from one machine to another in the peer to peer network till a match is found with key words. If there is sufficient match then the 'client' software says to the user: 'Someone wants to know about Microsoft, would you like to answer?' Now the user has total control over this situation. No one else knows that the user contacted has been asked about Microsoft unless he or she decides to answer the question. If the user decides not to answer the question then the software just goes around looking for other profiles. If the user does decide to answer he or she can opt to remain anonymous or give a name. Anonymity can be guaranteed because the answer is bounced randomly across all peer to peer computers so that all the person receiving the answer would know would be the last I.P. address that the system hit and not the originator. Of course in assessing how much an answer can be weighted in terms of accuracy knowing the answer originator is important, but he or she always has the option of encrypting so that only the relevant correspondent(s) will know.

The present system is called SHOCK (Social Harvesting Of Community Knowledge) and though at HP it is mainly used for informal question and answer there are a number of ways in which it can be useful. One is to monitor the kind of questions that are important to people and the other is to identify the gaps in an organisations knowledge by seeing which questions are asked more than others. But in general it could have wide application. In the health care sector it would be invaluable for contacting quickly a person who has some particular skill. The system is efficient, cheap and can be deployed very widely. In the world wide system used for IT security at Hewlett Packard, for example, a question asked in Singapore was answered by someone in Rio de Janeiro. At Stamford University there are many departments and often I might want to ask a question that I know that someone in the biology department could answer but I do not want to unnecessarily impose on that person by contacting them directly and making them feel obliged to spend time. SHOCK enables me to put the question so that people can reply at their convenience. We can also add an interesting layer by paying people according to the questions that they answer or not as the case may be. The level of question does not matter either. A child might want to know why the night sky is dark when there are so many stars about and some cosmologist or physicist somewhere in the world might like to answer.

Questioner 1 : Is there feedback in terms of how good the answer is?

Bernardo : We are working on a ranking system that would determine the worth of an answer. So far we haven't implemented it but in say, an area such as software development people do not like to spend their time fixing other peoples' problems unless they have to. If I have a problem I tend to just contact people who know a lot about certain things and then they might have to spend an hour with me explaining something. A system where people got paid according to how well they did would be a nice system. After all people who know how to answer question better than others should presumably get greater compensation.

Questioner 2: Are you familiar with the system that the World Bank tried to implement. Again the problem was one of incentive. The World bank has 10,000 people, 6000 of which are PhD's. These are supported by the other 4000 who are worse than the academics for withholding information. You ask them a question and they will say 'wait until I publish'. And the point about the World Bank is that it is important to share fundamental economic information about countries. In order to overcome the problem they tried to bring in a peer to peer system in which if an individual is asked a question, their first priority is to answer that question.

Bernardo: But how do you make sure that they do go with that priority?

Questioner 2 : Well that is a huge problem and the other problem is capturing the question for the next time so it not repeated.

Questioner 3 : In the Second World war the most efficient army was the German one, and the reason why that was so was that the people at the bottom of the 'food chain' trusted the people further up the 'food chain' (and vice versa). So if, for example, a commander wanted to know where his troops were, he trusted the people below. The people below were working on mission orders and so had a general idea of what they had to do, but they were trained such that in any kind of unforeseen circumstance they could adapt faster than the Americans or the British because of trust. Now what concerns me slightly is that, OK we're working in a capitalist society, but what (seems to be the implication) of what you're saying is that you can't trust the people further up the food chain because if you're down at the bottom and the person further up is earning more money why should you give them information. So in a way the mistrust is hardwired into the system (by the money incentive) and I'm not too sure how I feel about that, because the way that it's been resolved in the military situation is that everybody goes in at the bottom. The analogy at the World Bank would be that even if you're a PhD you start at the bottom with everyone else and that's where you're tested and chosen.

Bernardo: OK. You're bringing several issues here; one is of trust and in particular the reluctance of some individuals at some level in the organisation to reveal information to the top. I don't believe that the only way you can do this is by legislating. Perhaps you can do it the way the Germans do it but that may not work very well in the US. Perhaps in the army it's different, but in general you need to create a set of incentives. Obedience is something that you train soldiers to have as a reflex. It's a different situation when you have an organisation, in particular an academic one or semi-academic one like a research laboratory. And companies as well. You need a different kind of incentive, and not so much one in which you're punished if you don't reveal the information because you're ordered to do so.

There's an interesting problem to do with information bias. In a big company like Hewlett Packard we have lots and lots of business operations, but one kind is the sale of very large items, big super computers and so on and it is important for the managers is to know how many computers are going to be sold by the end of the quarter. This is a problem because if you ask the salesmen, they always low ball the prediction because they have a quota which they agreed to, and it's always low because they know they can fulfil it, and above it they get a bonus. If they agree to a

particular quota and they sell less they know they're going to be in trouble. What we have to do here is to align the incentives according to the kind of people we have in the organisation. One way to get information about how likely the sale of big items is, and you can only do this for large items, is to design a menu of different fixed income/bonus options and let the sales people choose. What we found was that the larger the bonus chosen, the more likely it was that a sale would take place. The salesman doesn't know this and in fact the relationship is a complicated non linear function, but the process is simple. At the beginning of a quarter they are asked to choose the form of their compensation. If a salesman doesn't think that anything is going to be sold next quarter he'll go for the fixed compensation. On the other hand if he knows that he's going to sell two or three of these big super computers this quarter he'll go for the big commission. So by just allowing the sales people to choose how they want to be paid you can extract information about the likelihood that a sale will take place.

Questioner 3 : Sure, I'm not making an argument. I can see how it works, but my contention would be that when the system doesn't work, when there's an error and the company loses money, who gets the blame? Because in the military system what happens is that the responsibility is on the person further up the food chain and they take the blame. In large organisations it seems that people spend a lot of time sending e-mails to one another, not because they want to transfer information, but because they are covering their butts so later on when something is wrong they can say, 'Look I e-mailed it to you, there is my cover.' And again comparing the military approach, paper is not used. Verbal commands or verbal agreements take place. I would be interested to see both systems in operation and I suspect, to be honest, that this system (money incentive) works a lot better than the military one, but what worries me is the underlying assumption that you can't trust your employees and that's why you have to stress the privacy because it's re-enforcing that ethos (of mistrust).

Bernardo: Well, it’s a legitimate question. We evaluate these things by testing them in the laboratory. We do experimental economics which was one of the things that I started when I came to HP and many of the methods I will describe later are tested in realistic situations. In the psychology laboratory, when you participate in something they give you a certain amount of money just to show up. The difference in my laboratory is, yes, you may get $25 to show up and you can make a lot more money if you do well, but you can lose the money which you've gained if you don't.

We call the salesman options that I was talking about earlier, where people can choose according to (what they see as) the likelihood of a certain event, ‘contingent contracts’. And we notice that people are very rational in that they eventually go and choose according to the optimum. It 's not so much a case of avoiding blame, which happens everywhere, but essentially putting in place a mechanism that induces the revelation of true information. How that information is used is a different story.

Questioner 4: I come from the industry myself and I want to ask you whether you get the best answers or answers from the people with the most time?

Bernardo You get both and we now have a system which does two things: it collects all the answers and assesses the frequency of each kind. The questions that are ranked highly are the ones that remain in the system so that people can get automatic

responses and we are now implementing a system which assigns stars to the best answers. This is a good system to deploy in say, a branch of the government or a very large global corporation, though we have been testing this in the laboratories where nearly everybody knows everybody else. The quality of the answers you get is a serious consideration, but it is what happens in a scientific or any professional community. In science people may have all sorts of opinions about the origins of certain kinds of phenomena, but eventually credit is given to those who happen to be right most of the time. It's a system that works fairly well in the academic world and technical communities because that's the way reputations are built. The community judges.

Questioner 5 : How do you handle context in this system because the value of an answer will vary strongly according to the context in which it was produced.

Bernardo: Well OK, at a general level your question is right and we don't know the context, but the questions that are being pulsed in this system are very specific. 'Does anyone know how to fix a JAVA code bug' for example, or 'does anyone know if there are cheap tickets to a theatre in New York City in November?' These are the kind of questions we see most often in the traffic.

Questioner 5: Is that the limit of what this system offers? I am thinking of many other issues of great interest in communities of practice which would go far beyond.

Bernardo: Yes, this is basically a question of asking: 'Does anyone know the answer to this problem?' That's the way it's being used. Now eventually it will evolve into something much more abstract But this is the way it's being used for now.

Questioner 6: How does this differ from software systems that are already being used in large corporations, where they trawl the Internet to find out what is happening? This is particularly true of multinational organisations with subsidiaries in different parts of the world. When the C.E.O. of a multinational company talks about something he or she knows about things that may be happening in other parts of the world. And not only is all the information about their organisation available from one source but they also have all the information on the competition gathered and stored.

Bernardo: Well they might know what has been stored or made public but many many times there is a part of the organisation that knows something and they have not been questioned. You cannot span a whole organisation and send an e-mail to 100000 employees asking them to answer a question. Everybody would be flooded with questions everyday from random and unconnected people in the company. It would be a very bad system.

Now usually the way it happens in companies is that the C.E.O. decides to give a speech on something and asks her underlings to find out about this or that. I can give you concrete examples. Every once in a while, even within HP it seems that some customer says 'By the way do you sell this?' And the salesperson says, 'Yes I think we do'. And the customer says, 'Well whom do I contact?' And the salesperson doesn't know and they start sending e-mails to the whole corporation and most people don't know the context of the message and they don't answer it. With a system like SHOCK it will go to the people who deal with that kind of knowledge, the sales people or the

manufacturing people. The system channels the information to people whose profile reflects it.

Questioner 6: I have a question on the privacy issue. As an employee would you have any choice about whether you wanted to be monitored and profiled in the first place?

Bernardo : You don't have to use the system for anything but work issues. Profiles are not available to anyone. Now you might say, but the infrastructure people can break into my machine. Well if you worry about that you shouldn't even use the system to find out about the weather because in principle that's not what you should be doing when you're at work.

Questioner 6: But I have a choice about whether or not to use the web.

Bernado : You have a choice about whether to answer the question or whether to participate in SHOCK and we do not force people to be part of the system. We ask you whether you would like to join. If you see that the potential benefits offset the possible liabilities then you'll join. If you think you'll get caught trafficking in pornography then that's something you have to decide. I mean, people say 'What if they know?' Well even without SHOCK these things can be found out. And there have been a lot of cases in the US where people have been demoted or deposed for keeping pornographic material.

Questioner 7: Can I look at my own profile?

Bernardo : Yes, and in a way that's a problem. I looked at my profile and I was horrified. You can modify your profile by the way and as a matter of fact we also offer a menu of key words that enables you to draw traffic towards you. So yes, you can see your profile which I did and here's a little anecdote. The profile is just a bag of words scrolling in front of you, but when I typed something, a horrible word, there were matches all over the place. When I typed extreme violence against other human beings there were again matches all over the place. What was happening was that, well suppose you look at CNN.COM or the BBC and you're looking at a web page that is telling you about terrorism or what ever. You might only look at the first paragraph of that page but the text below might have a horribly detailed description of some crime that was committed and it all gets pulled into your profile. I discovered the words 'mutilate' and 'mutilating' all over my profile but the system administrator had an even worse one so for the sake of not getting people too scared we don't encourage them.

So let me tell you about a different problem. Here's an extreme example. Every once in a while at Stamford there is an epidemiological study which involves a survey of the sexual mores of the students. The survey is circulated throughout the campus and most people respond because they remain anonymous. The problem for the researcher is that once he or she discovers that say a certain percentage of the female population has admitted to certain kinds of sexual practice, the researcher would like to know whether their partner was the same sex or the opposite sex. But that question cannot be asked because the respondents have remained anonymous, and that highlights the problem of privacy. If you allow people to identify themselves so that they can be contacted later then people feel they are revealing information about

themselves that might be used in all sorts of detrimental ways. In the UK in particular there are strong and stringent laws protecting people from data that is collected by their doctors and this means that epidemiologists are essentially devoid of any data to work with. The downside of this is that epidemiologists are unable to discover certain behavioural patterns or certain genetic characteristics that mean a propensity towards certain diseases and so on. And though the context in which I did my work is organisational it demonstrates that you can contribute to a public policy debate by bringing in technological solutions. The dilemma that the UK and the European Union have, is that if you have strict privacy laws you do not have access to data that would be useful to say pharmaceutical companies in developing new drugs. On the other hand if the data is freely available the potential for abuse is large. Under the strict laws in Europe people have to authorise the release of any personal data. Imagine say contacting 2 million people to get their authorisation and you get some idea of the problem.

One solution is for respondents to only be contactable through a trusted third party perhaps through some national agency. In the United States some people don't think the government is their friend but if I talk to people in Switzerland and Denmark they do. In the United States 30 years ago if a woman gave birth to a child she didn't want and wanted to have an abortion or give it away there were laws that protected the anonymity of that woman. Any many women counted on that law. But more recently the children pressured their own legislators to change the law so that they could discover who their natural parents were. The law was then changed 10 years ago and today any person over the age of 21 can ask for the identity of the parent. It is not that the government is malicious but simply responds to the changing pressures of society. So trusted parties do not really exist and the Internet is even worse because the data banks can be broken into by some clever hacker.

So one potential solution that I developed with Cat Hogg was to use an encrypting system which was published in *Nature Biotechnology* last year. In the survey the respondents remained anonymous, but we were able to contact them repeatedly. Now most of you are not interested in the cryptography but when I wrote the paper for Nature I had to explain it so I'm going to explain it to you with a physical metaphor.

The questions are posted on a public bulletin board. You use your computer to go to the website which then asks you whether you have a certain genetic marker or perhaps something like 'Do you like your manager?'. Now most people would like to answer the latter question anonymously, but I as the researcher, I would like to contact those who said 'no', perhaps to find out whether they disliked their manager because she was unfair or had bad breath or whatever. Let's assume for simplicity that the answer can be 'yes' or 'no'.

This is how it works. When you answer a question you effectively give me two identical unlocked boxes; one says 'yes' and the other says 'no'. You have the key that corresponds to your answer though I not know what your answer is. When I need to communicate with those that said 'yes' all I have to do is put the message in the 'yes' box, lock it, and so those people that have the 'yes' key can open that box and read the message, but the ones that said 'no' cannot. The system is designed so that it is not possible for someone to have both keys. Notice that this does two things (a) those people who said 'yes' can read the message and (b) you can put in the box that all the people who said 'yes' opened a key so that they can communicate with each other without the rest of the people knowing. This is how it works if you want to cover a

group of people who have particular characteristics: diseases, tastes, hatreds or whatever and there is no trusted third party as the keys are issued from within the system. From the point of view of the user interface it is very nice. If you answered 'yes' and the communication is for the 'no's then all you see on the screen are random characters. If you have the key because you answered 'no' then all you have to do is press one button on the keyboard and the message appears, it gets decrypted in front of you.

In order to prevent people having both keys we exploited a strange property of encryption systems. The way the system works is that there is a public key that is available to everybody and a private key that is your own which enables you to get into the public system. If you try to generate two public keys and two private keys and multiply the two of them to a very large prime number there is a theorem that says you cannot do that. You can produce two public keys that multiply to the number but only one private key. It is mathematically impossible to produce two public keys that multiply to a large prime number and also have two private keys. That’s a feature of the system.

But now you may ask for example, the question, 'What prevents me finding out who the Nazis are in a particular organisation by pretending to be a Nazi and asking the question: 'Who thinks Hitler was a great man?' Well there is nothing in the system as in real life that prevent you from doing that. But spoofing or the ability to pretend you are part of something when you are not is very costly so for the kind of data that we have in mind it is not going to be a big issue. The reason I say it is costly is that I read a book a year ago about how animals spoof each other. For example, when gazelles are gathered and hey see lions coming they start jumping up and down very high and this gives a signal to the lion that they are fit and the lions will not go after those that can jump high. As they get older so it becomes increasingly harder to spoof by jumping high. So you see that the higher the risk the greater the spoofing. If you really want to spoof a group of people into believing that you are something that you're not you have to put a lot of effort into it and the system is not really designed for that.

One of the disturbing things that I learned at an international conference in Copenhagen with many people from Scandinavian countries and Australia was that when they do surveys of young adolescents (13 - 17 yrs) which ask them about drug and sexual behaviour the children tend to be much more trusting if the answers are requested on computer than if they have to write them. This is because they have the mental model that a teacher or person in authority may recognise their handwriting whereas a computer appears totally anonymous. Computers are trusted more than people. And what is remarkable is that researchers report a 90% response rate. So it may be that when these people grow up they may live in a society in which they find the privacy concerns which you and I have are totally irrelevant.

When we designed our system and demonstrated that it worked people knew that the questions were originating from me, and we were a bit squeamish about what kind of survey to carry out. Fortunately two weeks ago we had a visit from the Harvard Medical school people who had set up an organisation called Harvard Partners with one of the largest medical centres in the United States. This is a centre for the genomics of diseases. These people are really going for the whole idea of personalised and individualised medicine where medication is targeted at the individual based on the personal gene characteristics and other personal information. They need to carry out this kind of survey perhaps at the rate of 10 or 15 times a day.

Questioner 8: Aren't you making spoofing much easier and cheaper to carry out?

Bernardo : Why? Because you only type things onto a keyboard?

Questioner 8 : Yes. It means that you don't have the trouble of getting involved socially. That's what makes intelligence gathering difficult.

Bernardo : OK, Let's assume that I'm part of a group that doesn't like to be discovered easily and you don't have to meet me face to face but can type something asking for personal characteristics. I will probably ask who is asking this question and may send you a message back saying 'What do you need it for?' Yes, OK, there's a lower cost but still in order to convince a suspicious group your behaviour is complicated and you have to appear to be a member of that group. Yes the electronic communication may enable you to hide your real person but spoofing is still costly.

Questioner 9: What happens if you lose the computer with the keys?

Bernardo : Well, nothing. You are the only person that has the password to get access to your own key. When you generate the keys you have to put a password in.

Questioner 10 : Is there is an apparent limitation in that your answers can only be binary, that's to say, 'yes' or 'no', 'black' or 'white' but not 'grey'.

Bernardo ; No, if you're asking the question: 'What if I have many answers?', the system has been extended to very large spectra of answers, but if the question or the answer is grey (in a semantic sense) then the survey is at fault. You avoid grey answers by the design of the survey, which is quite an art.

There is also a system using cell phones which finds people with particular characteristics. Suppose you want a recommendation for a movie you would like to rent. It's easy if the movie is about business but not so if it's some movie you and your wife want to watch and not everybody should know about it. What do we usually do in those circumstances? We talk to friends. People who have similar preferences to ours. Or it may be a restaurant. Do you take a recommendation just from someone in the street, or from someone you know who has the same preferences, which increases the likelihood of you having a satisfactory meal? We rely on asking people we know have similar preferences when we are dealing with sensitive or personal issues. How do I find out if someone has the same preferences as myself without revealing my own? When I was on a sabbatical in Paris what I discovered was that most people, and academic people in particular, are very political and so you have to be careful about expressing your own opinions. When you go to a party you usually trust a third party, perhaps the host or hostess, who will say, 'go and talk to so and so he is interesting' which are code words for 'you and he will agree on something'. And when you start the conversation you will probe very gently around some issues and decide whether you want to reveal say, your political preferences. Social human beings are very good at detecting clues about whether or not they are talking to a person with the same preferences.

What we wanted to do was solve the problem without a trusted third party. So here's a little bit of mathematics and I apologise to those suffering from math phobia

but I want to show you how it works. The proverbial types in cryptography are called Alice and Bob. Lets assume that Alice has A and B, which could be two phone numbers, and Bob has A and D. They want to know whether they have any numbers in common. To find out, Alice generates a secret key (x), that allows her to encrypt her numbers. Bob generates his own secret key (y) and encrypts his numbers. The phone numbers and the encrypted numbers that result are all integers. Alice and Bob agree on a common prime number which acts as a modulus. Alice then sends her encrypted integers A^x and B^x to Bob who encrypts them as A^{xy} and AB^{xy}. Bob sends his encrypted integers to Alice who encrypts them as A^{yx} and D^{yx}. Alice and Bob then know that they have A in common.

Refer to slide 12/13/14

Questioner 11 : Although this system is secret in the sense that you describe, what is the motivation for identifying communities of like minded people? Is it more sales to people who want to find a product?

Bernardo :What is its use? If you talk to the Nokia people, or Ericsson people, or anyone involved in this new world of people contactable by mobiles the magic word is ‘communities’. Companies want to identify communities of people with similar preferences. There is an interesting book by Howard Wrangle called *Smart Mobs* which is about a phenomenon that occurs in Europe and Japan, of teenagers going around matching each other on preferences and having meetings. Companies like Nokia and Vodaphone tap into these communities. You can target advertising of particular products. That's the interest of the thing.

Questioner : 11 The success of targeted direct marketing depends crucially on the response rates and a lot of research is going into trying to identify communities of people that are more likely to respond.

Bernardo : Right. I would now like to give you a provocative thought for the second session. You notice that I spent a lot of time showing you mechanisms that were designed to preserve the privacy of people. The question is, 'How many people care about privacy and how much is it worth to them?' This is a question to which we do not know the answer. People talk about privacy and yet on occasion they reveal everything about themselves. There is a large number of cases. People are asked on an Internet website whether they care about their privacy and people say 'yes'. Then they are asked questions and people reply willingly, about a lot of (personal) things. I want to design an experiment where people are paid to reveal something very private, and where the amount of money paid is a measure of how private that thing is. The problem is finding questions that people universally will consider very private.

One suggestion by a colleague of mine, and coming from her I don't feel I have to apologise too much for it, but it is gender based. She said that women do not willingly reveal their weight. So we could design an experiment in which a number of women are asked their weight. And they can't cheat because we can weigh them afterwards. You see we have to be provocative about this because it has to be something that people don't want to reveal but they will do so for a lot of money. We haven't done this experiment yet and the only thing I have done is to design an auction which will induce truth revelation. I tell you this not just to titillate but to provoke you

into thinking that most of the things that we think of as private we are willing to reveal in exchange for something. I don't know about English television but in the States one of the most scandalous things on television in the afternoons are these shows where people stand up and say the most incredible things about their personal lives in exchange for the publicity they're getting. So we must conclude that people have a willingness to reveal awful things about themselves in exchange for something, be it fame or money or something else. So I don't know. I know in the UK here there all sorts of laws on privacy yet you have something we don't in the States which is TV cameras everywhere recording you. And people don't seem to mind this. Maybe it's just a case of getting used to something and not seeing too many abuses advertised.

Questioner 12 : Perhaps the question is: 'How many people know the extent to which we are being monitored?' Because if you don't know when say, you're having a chat on the Internet (that you're being monitored) then you just trust. We are in fact being monitored more than we know.

Bernardo : But are you aware that there are 1.5 million TV cameras in Britain?

Questioner 12: Yes, but that's what is visible. But there are a lot of things like the Electoral Roll in Britain. Up until recently there was only one which was then bought by marketing companies for direct marketing purposes. This has now changed to two electoral rolls and you can opt out of having your details sold to marketing companies. I think that's something that people were not aware of.

Questioner 13: It's a bit odd, talking about your name and address, when there is almost a complete record of everything you have done today on the TV cameras.

Questioner 12 : But you don't have a name to a face.

Questioner 13 : Well, the technology is there.

Bernardo : Can I ask you, does your behaviour in the street change?

Commentator 15: I just don't want to be exploited. I don't want information going out that other people may use to exploit me. I don't want loads of stuff coming through the post box, people wanting to lend me money or attempting to induce me to do something.

Bernardo : On the other hand if you needed money. An advertisement for a loan maybe. Then that wouldn't be such a bad idea. Marketing and advertising also contain information it's just that.....

Commentator 15 : I don't want a load of junk coming through my letter box.

Bernardo : Well that's because the information is irrelevant to you. This is what the gentleman was saying before. If you could actually target the information to those that actually want it. Systems like SHOCK would allow you to spend that way.

Questioner 16: You don't think if you look into the history of interaction between primitive societies there's always been this reciprocity thing where we exchange and I think all we've got here is a 21st century model of that. We consciously or subconsciously exchange. I don't mind people knowing about me if it gives me the security of society. And I think that even the exploitation by marketing is fine because I don't necessarily have to buy the product. I think you have to judge the thing from a community base. If you had a group of muggers in here then I think they would have a different conversation about the cameras in the street. You have to look at what their exchange of privacy for what they get back means to them. And that's the critical question: 'Where is the line drawn?'

When British Airways had an alliance with United States Airlines about a year ago there was quite a lot of inter cultural exchanges and training for staff to understand cultural differences and one of the dangers we were alerted to was that a Northern European culture such as ours and a North Eastern United States culture were essentially very similar except in areas of personal privacy. The Northern European culture jealously guards personal information and we only share it with strangers reluctantly and when trust has been established. Whereas a North Eastern American will freely exchange almost all private information except for things like personal medical history.

Second Session

Bernardo : I now want talk about something we all care about and that is the future. I chose this beautiful painting by Caravaggio for my website for several reasons. It's called 'The Fortune Teller' and the first reason is to remind us that we're not the first people in history to care about the future. This is a picture of a young man going to the fortune teller to find out about his love interest because he thinks the fortune teller knows something about the future. It costs him money. And if you look at the picture carefully you will see the fortune teller is not only telling him his fortune he is stealing his ring which is perhaps an interesting metaphor for what happens when you go to some consultants who tell you things that are going to happen and you pay them a lot of money and then they disappear. And this is a problem. People and institutions, governments and academic circles etc., spend an immense amount of money and energy trying to find out how to invest resources or what research to do.

I'm going to show you my favourite quotation of all times that comes from a great hero of mine, Neils Bohr the founder of quantum physics. He said it is hard to predict anything, especially the future. People are great at predicting what has happened and they're always saying 'It was obvious that it was going to be this way' or 'It was obvious that the stock market was going to collapse or the Internet bubble was bound to burst'. So it's a difficult problem and I think Neils Bohr was referring specifically to a piece of work that was shown to him by an economist who was able to show that his prediction was in line with the data up to that date.

I was involved years ago with a student of mine and a colleague at Stamford in trying to use neural nets to simulate how humans foretell the future by extrapolating from the past. We didn't do that well and found that there were better ways. What I want to show you here is something that depends on the kind of intuitions that people have, how we attempt to ensure that they tell us the truth and how we validate their predictions against some objective data.

First let me tell you how organisations usually predict. Typically they ask the consultants what is going to happen on sales of certain products or how the economy going to go and so on. Then they have lots of meetings and usually they try to reach some kind of consensus about what is about to happen, whether to invest dollars or Euros or pounds in this or that area, what is the future of certain kinds of activities and so on. What usually happens at these meetings is that some powerful opinion prevails and if that person, or persons are dominant, that's what gets broadcast apparently as the result of a democratic process.

But voting can be a phenomenal way of getting at the truth and it is the basic foundation of jury trials in the United States and other countries. The basis of democracy rests on a theorem by the Marquis de Condorcet who was asked by the aristocracy at the time of the French Revolution : 'What do we want, a totally wise and brilliant king or the opinion of the majority of people who are totally illiterate, uneducated and brutish?' Condorcet, even though a Marquis, had libertarian ideas and decided to look at the matter mathematically. What he showed was remarkable.

Suppose there is something which is either true or false. For example, I am accused of committing crime. Lawyers will argue both ways in order to convince the jury. If the probability is slightly greater than half, in other words analogous to flipping a slightly biased coin, then intuitions are a satisfactory basis for attaining the truth. Moreover, as the group gets larger the probability of attaining the truth approaches 1,(certainty). On the other hand if the probability is slightly less than half, the probability of being right is zero. This is important because it means that voting can in principle be used to make decisions providing you have sufficient information. However, compared with the methods I will show you voting is not a very good way of getting at the truth.

A couple of years ago a group of economists at Caltech and the University of Arizona researched a set of ideas that were actually articulated first by von Hayek and then Lucas in 'rational expectation theory'. This was based on the assumption the that markets are incredibly good at aggregating information because millions of people participate in the market and all the little titbits of information that come one way or the other, eventually get folded (incorporated) into one value, the price of an equity. So the principle is that, if you try to decide whether a company has a future, and you have to qualify future (say 6 months), all you have to do is to look at the future value of the stock that's traded today (i.e. futures). This does not involve insider trading. It's all these people getting all their pieces of information and folding them into their buying and selling. This is the idea behind the notion of the 'intelligent' market.

So the proposal was that outcomes could be predicted using the same technique as in markets. Instead of buying a stock in a company people would buy stock on an outcome. For example, suppose people buy stock on the proposal that it will rain tomorrow. According to the principle, if we just look at the value of the stock we can tell how probable it is. There's a good example of this in operation on the University of Iowa's Dept of Economics website. Many years ago they studied the information market to predict presidential elections in the United States and every subsequent year there was an election they got closer to the truth than the Gallop poll. If you look at the presidential election between Gore and Bush the accuracy was remarkable. Of course at the end, when the difference was only a few hundred votes the price was 50/50, but at the time my family and friends were surprised at my ability to predict things when all I was really doing was going to the Iowa market data. The Iowa market showed that when the system is very large the predictions concerning

presidential elections are very good. Another information market was based on Hollywood stars. People bought stock in Hollywood stars and by observing how shares were trading on popular names it was possible to predict the outcome of a movie in which they were going to participate.

Such markets work because they share some properties with real stock markets. In the stock market hundreds of thousands of people participate and therefore there is, firstly, a lot of information aggregation, and secondly, the market cannot easily be manipulated because there is a limit to the number of shares that can be bought. I wanted to try a similar operation in an organisation but problems arose when too few people participated. When we ran a market for predicting sales, the sales people were reluctant to participate and there was a problem of not enough cash circulating for the buying and selling of shares. Then there were 'information traps' caused by rumours which circulated and grew with the result that the majority of people ended up buying the wrong shares causing a 'bubble' and then a selling frenzy. On top of the difficulty of motivating a sufficient number of people those that did participate were easily manipulated. It was very easy for a small number of people to buy a lot of shares and manipulate the market.

We therefore tried a new approach with Kay Yutchen and Leslie Fine in which participants with good predictive talents were identified and peoples' 'risk attitudes' were assessed. The technique avoids the pitfalls of small groups, induces people to aim at the truth in making predictions and involves an interesting non-linear way of aggregating information. The original paper on this appeared at a conference of the ACM on e-commerce (2001), and we have a long review article that will be in this month's *Information System Frontiers*.

I gave you the example of predicting whether it is going to rain or snow tomorrow. If I say that I going to pay you if you are right and you will have to pay me if you are wrong you are likely to do a little research before you give your prediction. Now I realise that people can be very diverse in their attitude to risk. A very risk averse person may actually refuse to predict by hedging their bets. They will think 'Well, perhaps I'll say 50/50 because I might not get all the money, but I won't have to pay so much if I'm wrong'. On the other hand, a risk seeking person will tend to amplify the information and gamble all the money on something that other people might consider pretty uncertain. In real life we know people that are like that and we know how to re-normalise what they tell us to take such things into account. You cannot do that with people that you don't know. Most of the experiments that we do are with students from Stamford University and what we needed to do was to first assess the risk attitudes so that when we aggregate the information we could take it into account.

The mechanism I thought of takes principles from portfolio and finance theory. It regards the information in people's heads in the same way as assets and money is invested in a similar way. The exercise has two stages: first we run a very short and very imperfect information market in which we assess the participants attitude to risk and second we ask them to generate predictions on outcomes which are then aggregated.

Refer to slides 23/24

In the first stage we have about 13 or 14 students in the room at the same time and the whole process takes about 40 minutes. Each participant draws different

coloured marbles one at a time from an urn on a computer screen which contains 10 different colours one of which is more numerous than the others. They can buy three different kinds of security ticket: A, B and C. Each of these 'securities' pays out different amounts of money at the end of the session. As each participant draws a fresh marble and puts it back so their intuitions about which colour is the most numerous grows. At the same time they are allowed to buy and sell marbles of different colours using the 'securities'. The market runs as the stock market might and people do not necessarily trade according to what they think the end result will be. The other day I saw a student buying blue and I said 'So you think blue is it?' and he said 'No, but everyone thinks it's blue therefore I'm accumulating blue so that I can sell them before the end of the game'. You see, there's a bubble. But the object of this stage is for us to extract a risk attitude without the people knowing. At the end of the game we reveal the most frequent or winning marble and a participant gets paid according to how many they have of the winning colour. The exercise is synchronous in that all the students play at the same time and we assess each persons risk attitude by extracting a number which tells us how risk averse or how risk seeking they are. We find this is an amazingly constant number even if we ask the students to participate in other things.

Refer to slide 25

The second stage is easier in that there is no market and is asynchronous in that each student is asked to make a prediction in turn. There are again 10 colours and once again they draw colours which they observe. They are also given 100 tickets and have to spread the tickets on the colours according to how they assess the fractional probabilities. At the end of the session we reveal the true state of the system and they again get paid.

Predictions can then be aggregated according to the following formula, which is a version of Bayes formula and I apologise for some math, but I'll tell you what it means. Remember each person is making an assessment and we need to get a single prediction on the probability P of an event s conditioned on I.

Refer to slides 26/27

In the formula each Ps_1 , Ps_2 , Ps_3 etc. is the fraction of the tickets that went to each of the colours and each β_1 , β_2 , β_3 etc. is the behavioural risk rating of the player. If the player is risk averse β is a fraction above 1 and if risk seeking a fraction below 1. In fact we discovered risk neutral people don't exist. The β factors were extracted in the first stage of the experiment by seeing how much money a player has left compared with how much money he or she invested in the market. Because a player knows that there are 10 colours and that one will pay out and the total possible winnings it is obviously silly to invest more than can be recouped, but people who are very risk seeking tend to over invest and this lowers their β factor. One the other hand those left with marbles and cash have a higher than 1β factor . That ratio of securities against cash that you have at the end also tells us how risk averse you are and also since the game is played for 40 minutes each trading session refines the βs.

Questioner 17: Do people know you're doing this?

Bernardo : Yes we ask them to play the game.

Questioner 17: That you're trying to work out their behavioural risk factor?

Bernardo : No, we don't tell them that but we do ask them at the end if they thought the game was fair. We tell them we're going to measure their behaviour but we don't say how.

Questioner 18 : But isn't it unethical not to tell them?

Bernardo : We don't tell them how because then they'll gain and be strategic. And of course though peoples' risk attitude in this kind of activity is constant it doesn't mean that they are the same in science or the real stock market or love.

We do run a large number of experiments, but you must realise this is one which is totally under my control in that I determine the probabilities of the marbles in the urn and so I know outcome. Whether or not what the participants say is truth revealing we can consider later but there is an inducement to play well because they make money. Of course the amount we pay students would have to be increased if it was HP executives but there is a lot of empirical evidence that the amount of money doesn't really matter. Doubling it doesn't make any difference even though an hour of an executives time is worth a lot more than we pay. But it's often just the ability to say 'I've won and got a cheque'. So there's no real correlation between the truth and the amount of money at stake on the table.

Questioner : 18 What distribution of βs do you get?

Bernardo : We tend to see that there is a learning period. At the beginning of the game people tend to be a bit risk averse because they're not familiar but they quickly convert to something that seems constant. But we do see a spectrum of risk attitude across different groups. Sometimes with some groups the norm is towards the risk seeking and sometimes towards risk aversion. It's single peak, but then we do only have 13 or 14 people.

Questioner :19 Do you test the changes in risk attitude according to communities in which there is interaction between individuals?

Bernardo : Well they're not allowed to discuss things.

Questioner : 19 Well I am trying to make some connection between the experiments and the real world and the way people deal with things in the real world is not individual.

Bernardo : OK, I will show you later how we deal with the fact that people communicate with each other. In these experiments however, people don't communicate except in their pricing when they're buying or selling at the first stage.

Refer to slides 29 -38

So lets look at the results. The blue line is the distribution that we put into the urn. In fact if you were to speak like a probabilist this is what an omniscient being would know. Notice there are two peaks the main one at D and small one at I. So participant should find an increasing probability at the former but also a small probability arising at the latter. What we would like our participants to do is to ultimately map the whole distribution and we have a technique for measuring the difference between two distributions. The first graph shows the two distributions when no one has played the game and note the red line shows that all possibilities are equal.

What happens as you start aggregating the results? Each player has played the game 17 times and the second graph shows the aggregation for one player which isn't very good. Now if you have two players and you aggregate them according to the formula we can see no improvement. For three players and then four we can see a bubble starting at I. The group obviously thinks that this is the more probable event. Five players and we can see it getting more concentrated on I, but we can see a little blip coming at D. No real improvement for six players, but then for seven, eight and nine there is rapid improvement and eventually we have something close to the whole distribution. It didn't predict it exactly but it got the variance in the whole distribution. Voting would not give us this. It would enable us to see which of two alternatives is more probable and the probability would increases with the number of people but it wouldn't capture the whole picture.

Questioner 20 This is the result of the second stage?

Bernardo : Yes. They make extractions of the ball and then we say ‘place your bets. 100 tickets on all these colours’. Now suppose I had used the market in the first stage for predicting the distribution by seeing where the money is going and comparing this with the best person. this is given by the yellow line in slide 39. Using the market for prediction is not very good because even though it was on the right track at D there was a bubble around F when everybody was accumulating a lot of shares. In the market the variance in distribution is enormous whereas in the game it is sharply defined. If you talk to information theorists they will say that their theory indicates that a system cannot do better than the best player. But to me the notion that the group is better than the individual at predicting seems obvious. What we see here is that information, aggregated in a non linear fashion according to the formula, does better then the best individual.

Refer to slide 39

Questioner 21 : Sorry I didn’t follow how you use the market for prediction.

Bernardo : Where did the yellow go? Well we say ‘How much money on each colour?’ The Iowa market does well because there are 2 million people playing but when there is only thirteen it's not so good and it's an interesting problem in intelligence gathering and assessment that once you have the data you can all sit round the table and discuss it so the information becomes public. In real life we might draw little titbits of information from an urn but we also read the Times. So what do we do about this? When I make an assessment of the weather I'm already incorporating the knowledge that I picked up from the newspaper or radio and this would corrupt the

kind of experiment we've been considering. When I gave a talk in Stamford there was a big debate about what constitutes public information. Everybody gathers information which is public but what goes on in our heads is private. This may be a signal from the world that everybody saw but the way I manipulate it in my brain is different from someone else. In order to understand prediction in the real world we need to know what the public information is. It's impossible to take the bets in our heads and compare them with the bets in other people's heads. For a rational assessment need to subtract the public information.

Questioner 22 : Why do you want to eliminate the bias of public information? The real world has such information.

Bernardo : Yes but if you and I read the same article in the newspaper and we have to make an assessment of a probability, we have to know whether it is contaminating our assessment or not. And we would like to eliminate that.

Questioner 23 : In a jury system it would be very useful if you could eliminate it.

Bernardo : Yes, and I will show you what public information does.

Questioner 22 : But we all have non-linear functions operating on the information whether it's public or private.

Bernardo : I agree. We process it differently from each other. But what we did in one experiment was to have one draw which was public and when we ran the experiment we got public information bias. What was happening and it shows in the math, was that we got this peak at H which is an amplification of this public signal that everyone started giving an enormous amount of emphasis to, and didn't trust their own private information as much. This was a set back because you cannot go into an experiment where you're going to predict something about the real world if you don't know what is public and what is private. And we showed that public information always gets amplified.

Refer to slide 42

Questioner 23 : In this game you have private information because you know the colour of the ball?

Bernardo : Because you extract it yourself.

Questioner 23 : Right, but in real life what is private information? It's the people you talk to, your own thought processes. The information you gather around.

Questioner 24 : Not intuitive?

Bernardo : We don't ask people how they do this, but the issue is : if we have ten technologies and we want to predict which one is the more probable in say a years time and everyone has been reading the Economist which says nanotechnology is the greatest thing, that is going to be stored.

Questioner 23 : But there are many sources and what we call public information is not (easily definable?) and I would expect that the best we can do is assume we have an open society and information is distributed in a good fashion and the public information will be a good predictor.

Benardo : Well I'm not talking about predicting events in general. This is also a real thing. When we ask people to predict revenues we basically give them four options: is it going to be 10% higher or lower than last time or 20 % etc. We give them discrete events. They play the game with the marbles and we get the risk attitudes and then on the basis of the information they get they have to place their tickets. And at the end of the month they get paid. If we're going to talk about general things it's a different story. This experiment may be circumscribed in that way but it's not that contrived.

The solution to this (public information) problem actually came from Kay Yuchen who remembered a game he saw on television in Japan. They bring couples onto the TV set and put them in separate cubicles and ask them questions and if both of them answer the question in the same way they get paid a lot of money. We have an adaptation called 'The Matching Game' in which two people play. Suppose, and this is a gender stereotyping, but suppose that the man wants to go to the ball game and the woman wants to go to the opera. If they both decide to go to the ball game he gets three tickets and she gets one. If she chooses opera and he chooses the ball game no one goes anywhere.

Refer to slides 43/44

The idea behind this is, if you can guess what the other person is going to say or likes or going to bet, you will make money. So we have a further stage in the experiment which lasts about 15 minutes where we have people are taking the balls out and some of the information is public and some of it is private. We then ask individuals to make an assessment of what the others think about certain frequencies of balls. They have to guess what is public and we can then correct for public information.

Refer to slide 45

So this is an example of the matching game where even though the people know nothing about how the balls are appearing to others they make a guess of what is public. They don't know what is really public, since they only see their own screen. They see their computer extract a ball , but they don't know if it is being shown to all computers. We allow the whole room to see a certain fraction and other extractions are private.

Questioner 23 : There's no communication between them?

Bernardo : None, otherwise there would be collusions.

This allows us now to move into the real world. Yes, it is a circumscribed experiment where people do not have an opinion and we ask them to have very discrete choices. The experiment is designed in such a way that we can divide the events that we want to predict into discrete ones. We're now doing a revenue

prediction on a monthly basis and we have 13 executives and many volunteers who are not executives. Every week they get together and go through an assessment, revise the probability and make a prediction. At the end of the month we see how it turns out.
We were going to do this with a different division that sells very big equipment, but its quarterly revenues were too small and people are not interested if the payoff is not for 6 or 7 months.

People in government are very interested in this and we are confident that this kind of system would work very well. However I was involved in an experiment three years ago that was actually a failure from a social point of view. I had the idea of getting every member of the laboratory at Xerox Park to place bets on outcomes of projects that people were presenting to the management. People became incredibly upset when they saw how their colleagues were rating them and it was terminated by management. That was a case of not paying sufficient attention to the social aspects of the thing. But we do plan to do a series in terms of technology assessments, revenues and so on and we want to announce the winners because that's a good incentive for people to try hard. We have now changed the entire population of the people in the experimental pool and have students from other colleges.

Questioner 24 : In real life I guess you must have very specific issues if you want to get predictive results. Can you think of any other examples? I ask because several years ago I talked to someone calling himself a futureologist and they claimed that through constant feedback of results to the group of experts that they had a very high degree of accuracy. This is different from what you are doing here. But it needed very specific questions or issues

Bernaro : But the feed-back was on the real outcome?

Questioner 24 : No, the feed-back was the opinion of the group at each pass.

Bernado : There is a system. Delphi. Rand Corporation ran this kind of thing in the 1950's. I talked to some of the people involved and it was not very good. I think it is important to have simple and specific ways of presenting this forecasting exercise so people know exactly what they are betting on and what the outcomes will be. When things get fuzzier it's harder to assess the results and our experiments are about making discrete events specific and possible so that people can bet on them. There are certain things all of us can predict because they are so periodic; like it's colder in winter than in summer but that's not so interesting.

Questioner 25 : In the sales forecast how do you take care of the public information?

Bernardo : By playing the matching game. There are events and each person makes a prediction on each event and we ask 'What of these events do you think other people know about and have a high probability?' They don't asses things perfectly but when there is money riding on the thing they do well enough.

Questioner 26 : I'd like to distinguish two kinds of predicting the future and I think you're talking about only one of them : One is predicting the future where the group of people making the prediction affect the future and one is where they do not. The

one where they do is the stock market and I don't believe what you are talking about applies there. Certainly I believe that that is very much determined by public information. The second one is where the events will occur independently of what the people advising you do. I think this is what you are talking about what you're doing is testing or trying to optimise the sample in order to give the best predictivity. In doing so you are essentially measuring the non linear functions that they are using in order to ascertain to what extent they give good results and trying to optimise your set to effectively determine the controlling properties of the system and only then asking them to predict how those things will evolve in the future given the facts of the past.

Bernardo : Correct, and as a matter of fact that brings in a big issue. If you have someone who can affect the outcome in the market then it's disaster.

Questioner 27 : Aren't most markets determined by the speculators that operate them?

Bernardo Yes, and we're talking of a group of thirteen or fourteen people. Supposing you're a sales manager and your group has decided that sales will not be as good as they last were and there is money is on it. You might as well tell them to take a vacation because the people will make it happen. This is the moral hazard problem of having those that affect the decision participate in the market. Markets can be controlled and participant's decisions become strategic.

Questioner 28 : Another comparison with the real world will be: when you're drawing balls from the bag there is no doubt that the ball that you're given has a particular state or not (colour). In the real world large organisations will have prior beliefs about what the state of affairs is and when a new piece of information comes into that domain, and runs counter to a belief there is an overwhelming predisposition to discard the information as being false even though it may be true and therefore signals of change tend to be disregarded until far too late. Whereas in the experiment we've got, there is no ability to disregard the data.

Bernardo : Yes, you disregard it at your own peril, but at HP we had a division that was doing very poorly and making quarterly forecasts and even though they had all this data coming in the forecasts were low balled. In these experiments we're taking people from all sorts of ranks and they have to make an assessment among a number of discrete events. You're right, and the filter is something that can direct them the wrong way, but their overriding need to make a profit in this game may override the mechanisms you are talking about. We can find that out empirically.

Questioner 28 : How do you select the states that you are going to present to the players? This is one question. And the other is what if I come in (to another session) and think in a totally different state?

Bernardo : How? Every week you come and play in the game and I ask you at your desk to tell me out of the four or five possible outcomes, i.e. are sales going to be 10%, 20%, 30% or 50% higher or lower than last quarter and I give you 100 tickets and 10 states and you have to place a bet.

Questioner 28 : Well OK, but the next question is: will you not have different outcomes depending on the particular participants that you have in a group?

Bernardo : Oh, I agree, and we would very much like to measure how accurately a group does depending on it's composition. This is part of the game's structure.

Questioner 29 : Have you made a comparison between American and Europeans?

Bernardo : What is your intuition for thinking they would do differently?

Questioner 29 : Well different cultural dependence

Questioner 28 : I would expect different kinds of wisdom in an organisation. I would expect management to have some sort of equivalent about who to bring into a (sales) meeting of this kind. I mention that because once you define the problem as: 'are we going to sell 10% or 20% or 30% etc. I'm just trying to connect this to the real world of organisations

Bernardo : So far my understanding of it is that it helps in short term forecasting with well defined situations.

Questioner 30 : Does this mean you can get rid of managers?

Bernardo : Well, managers are not just used for prediction, but this is an information aggregation tool that allows us to construct a probability of future outcomes with a short horizon.

Questioner 28 : If you connect this to organisational theory it may be possible to establish where the valuable information is for a particular problem within the organisation and you could build up links or connections. People are all the time bringing different bits of information and the interesting thing for an organisation is to pick at all these segments and to balance them according to different purposes and interests,. So this is an interesting tool if you can go beyond the boundaries.

Questioner 31 : I want to take this one step further because I think you said that the past informs a decision for the future (extrapolation). In a turbulent and changing environment the emergent patterns cannot be predicted from the past and therefore managers in many organisations today need to be able to adapt very quickly to totally new situations and this is our greatest problem in terms of prediction.

Questioner 28 : Well that's a good question and people do talk in an exponential way (about IT ?) without any idea of what the technology will be

Bernardo : I don't disagree with what people are saying, in that there are situations where you cannot predict from the past. On the other hand you don't hire the first person off the street to be your manager. You don't say 'since we don't know anything about the past I'll ask anyone'. The past might not be very useful but the experience of a person who knows how to filter information is always someone you need. And this has applied throughout history. There was no time when things were so linear. I think

people who are able to make these predictions are those who have learned to read signals that other people cannot. And they do well.

Questioner 32 : I'm fascinated with this idea of people being different as regards their willingness to take risks. Is that really transferable from marbles to other activities?

Bernardo : Well that's much more cognitive psychology problem. We are trying to see how far we can go in having a different model from the market that was played to get the risk behaviour in other things they have to predict. So far we have kept things in the same vein. We do the first stage and then present people with situations that are the same as the ones with the marbles.

We have had some very circumstantial evidence that these risk attitudes tend to be fairly broad and extend over many other games. We are now running an experiment on cheating and we see the same kind of behaviour. People are allowed to cheat and make money but their reputation may suffer. Risk taking people tend to be the ones that will (cheat) more often.

Questioner 32 : Would SHOCK (help in this?).

Bernardo Yes, we want to use SHOCK to identify people for the experiments, though it's hard to convince highly paid vice presidents to play. But the groups are very diverse. People from sales, accounting, and there are some very senior ones.

Questioner 33 : Don't these simple experiments offer us an interesting sociological exercise? A way to look at people with the same background but different cultures such as Swedes in Sweden or Swedes in Minnesota or Swedes in some other culture to test whether on this simple model you have the same distribution of risk aversion in these different groups. That would test the relevance of nature as compared with nurture.

Bernardo : That is a very nice suggestion.