



25 APRIL 2008

MATHEMATICS OF CURRENCY AND FOREIGN EXCHANGE

NORMAN BIGGS

I have to say that when I set out to prepare this talk, the material that I'm going to actually present today I imagined would be approximately a quarter of what I was going to talk about, but I got into this subject and I found that I was really rather satisfied with the conventional wisdom on some of these things, so what I'm going to talk about is - it says there - arithmetic at the end of the 13th Century, so this is definitely medieval, and I apologise for not going into the later developments in foreign exchange, which have a great deal of mathematical interest, but in a different way.

So the talk's actually going to be in three parts, as usual. I'm going to start by talking a little bit about the financial, commercial background; I then want to talk a little bit about the technology, as it would be said, that is the arithmetical tools that were available; and then I want to see what you can deduce by putting them together.

So we begin with some money. If you were a Viking, at the beginning of the 10th Century, this is what you would have understood by money. This is the Cuerdale Hoard, which was found - Cuerdale is in Lancashire, and all the objects are bits of silver - we now call them hacksilver - and among them, there are some rather special bits of silver which we call coins. They were not of course Viking coins; they were Saxon coins, Anglo-Saxon coins. The Vikings took them over and treated them as part of their money system.

So what do we mean by money? For our purposes today, 19th Century analysis will suffice. A modern economist would give a series of lectures at this point on the functions of money and its uses, but Jevens wrote a book in 1972, which was very popular and very influential, and he found it worthwhile to distinguish from the beginning between a medium of exchange, by which we mean the money objects, the coins the hacksilver, and so forth, and the accounting units which are used to measure value. He believed that money arose mainly because these devices made barter, or avoided the inconvenience of barter. Modern economists would think of other uses for money, and many other kinds of money of course, but let's just think about barter for a moment.

Here we have a picture from the 11th Century, which was printed by Salzman under the title of "A Simple Bargain". So, we think we know what's supposed to be happening here. There's a chap who appears to have a chicken, and a chap who appears to have something else, and they're trying to barter, but it's not as simple as that. First of all, the something else looks a bit like an old sock, which wouldn't actually be much good for barter, and then, if you look a little more closely, you'll see that the chicken man has in his hand a coin, in other words money, but it's not clear whether he's received the money from the sock man or whether he's actually giving the money to the sock man. I've looked at the original of this, or at least what Salzman says about it, and it's still totally unclear what this is supposed to represent. Those of you who get the BHSM bulletin will have read a paper by John, from the Open University, John Mason, who explains how barter came to cover an enormously complex operation in the later Middle Ages, with the intervention of money in all forms, money and credit, whether it was ready money or not. So barter wasn't all that simple.

Coins, however, did simplify the situation, but in order to use coins in trade, you've got to know certain things about them. In medieval times, a coin was worth, that it was accounted for, in the terms of the precious metal that it contained, and you could only measure that by actually doing two things. You first of all had to ascertain the mass, which you could do by weighing, but then you had to work out the fineness, in other words, what proportion of the mass was actually the precious metal. That second thing, the assaying, or finding the fineness, was rather difficult. There are methods known, and have been known since antiquity, for assaying both silver and gold, but they are not something you can do on the spur of the moment. There is a touchstone method, but it's extremely unreliable.

So weighing was normal in trade, and here is a rather depressing picture. These are some poor slaves are being sold, but you can see that the essence of the transaction is that the coins that are being paid for the slaves have to be weighed, and so they are not accounted for by the number of the coins - it is the actual amount of silver or gold that is concerned.

Some improvement came along in Western Europe, Christian Europe, when gold coins were reintroduced. There had been gold coinage in Roman times, and it was reintroduced, gold coinage, in Islam of course, where gold was in supply from Africa, but gold coinage didn't come into Western Europe until, wasn't reintroduced until the 12th, 13th Centuries.

So here we have the four representative coins from that time: starting off with the typical Islamic dinar; and then a coin, Castilian coin, imitating the dinar, the second one, obviously the same design, but with a different kind of lettering; and then we get into the more typical European-type coins, something called an Augustale of Frederick II; and then, playing a large part in the rest of this talk will be the florin, which was the Florentine coin which began to be minted in 1252.

Because there were a lot of these things coming into circulation, it was necessary to know the fineness of the coins. I've mentioned the weight. We assume that a merchant could ascertain the weight of an individual coin, but in order to ascertain its fineness, he either had to do a very complicated operation, which was beyond the bounds of practicality, or refer to some reference. Here we have the list made by Pegolotti, a Florentine merchant, which was compiled around 1300, and these are just some of the gold coins that were in circulation at that time. I've put little marks against the four that were actually depicted on the previous slide.

So the first one is the florin, and the Florentines claimed that this was 24 carats fine, so you see something there which says the number 24 is the crucial one - 24 meant pure, so everything was divided into 24 parts, and instead of a percentage, which of course they didn't have, it was how many parts out of 24 were gold. The Florentines believed that their florin was pure gold. If you look down, you'll see that the Castilian coin was $23\frac{3}{4}$ carats fine; and the Alessandrian bisant, the Islamic coin, was believed to be $23\frac{1}{2}$; and the poor old Augustale, or Agostantini as it's called here, was only 20 and something, 20 and a bit, carats fine. In a transaction of course, large sums of these things were being used, and it was necessary to know this fineness.

So that's the starting point. Now we move on to the international situation.

So, if you want to send your wool to Florence from London, then, well, how are the Florentines going to pay for it? They could send florins, but that would be a risky business, sending a large amount of money by ship, and so in the 13th Century, the mechanism of the bill of exchange grew up. I'm going to just very briefly explain how that worked. This is based on the account given by the eminent medieval historian, Peter Spufford, who has made a study of all these subjects.

So this is the situation for sending a shipload of wool from London to Florence, and how it's going to be paid for using the bill of exchange mechanism. So there are the four parties, two of them in London and two of them in Florence. On the left-hand side, they're the principals if you like, and on the right-hand side, there are the bankers. What happens is that the agent in London sends the shipload of wool to Florence, and various entries are made in the books to account for that. The remitter, as it's called, in Florence then pays for the wool, but pays a banker, a drawer, in Florence, in Florentine currency, florins, and various accounting things are made. In order to settle up, we then draw up, the banker draws up the bill of exchange, which he sends back to the remitter, and also notifies an agent of his, the payer, in London, and the bill then goes back round the system, is presented by the payee in the London to the payer, and the payer then pays the payee, this time in London money, sterling. So the point is that, in order to carry out this transaction accurately, there is obviously going to be some mathematics involved, and the next step is to ask what technology, that is what mathematical tools, arithmetical tools, were available at this time in order to carry out such transactions.

Now, I'm glad to say that Fenny Smith and I agree on several of the points I'm about to make, so excuse me if I'm repeating things that she said. But we all know that you can't do sums with Roman numerals, and of course the Romans weren't stupid - they managed to manage their empire, large empire, for several hundred years without actually trying to do sums with Roman numerals. For sums, they used the abacus, the calculus, the counters on a grid, and they became very adept at doing that. The word "abacus" can be confusing because many people think of abacus as meaning some eastern object - sliding beads and so forth. Abacus, in this context, means any grid or plan onto which the counters could be used. The other problem with it is that of course the abacus would quite often be just scratched in the sand or drawn on a slate, and so very few of these things have survived, and the counters themselves, although there were undoubtedly lots of the counters, when they're dug up nowadays, the archaeologists tend to say that they were gaming counters or things of that kind, that they're quite often mis-classified.

There were other ways of doing arithmetic, and here's a particularly horrid one, moving into the medieval period now. This comes from Bede, and this is finger reckoning. Although it's great fun, I cannot really believe that it was a very practical method of doing arithmetic, and it's really just there for interest.

However, when we come to the turn of the millennium, things begin to improve. This is usually associated with the man who became Pope Sylvester II, who I shall pronounce as Gerbert - if anybody knows better, please tell me - who is credited with various steps in the improvement of arithmetic. It's confusing, because not only is he credited with introducing an improved form of abacus, but he's also associated with the introduction of the Hindu-Arabic numerals. So let's see what actually was available, and for this, I'm going to rely on two accounts, and I rely on these two because they are basically confirmatory - they both say roughly the same thing, although they don't appear to a great extent in the published accounts in the history of mathematics.

The first one, this is a book by Florence Yeldham, published in 1926, called "The Story of Reckoning in the Middle Ages", which I would recommend. She gives this manuscript from Ramsey Abbey, and she describes the accompanying instructions from the manuscript. This is, if you like, an improved form of abacus. The first thing to notice is you've got the Roman numerals heading the columns, so you've got...working from the left, you've got units, tens and hundreds, and then repeated again, units, tens and hundreds in the thousands, if you like, and so on up to...I think there are nine lots of the columns. So that's Roman numerals. But then, hidden amongst the arches at the top, you will see Hindu-Arabic numerals, and also, while we're talking about the writing, what's written along the bottom are in fact the Roman names for fractions, which indicates that division sums were being done here as well. So, the other important thing, which is not apparent from this diagram, but which is confirmed by the descriptions, is that the counters that were used in Gerbert's abacus did in fact have the Hindu-Arabic numerals written upon them, and the calculations were done by moving these numbered counters around.

A similar account is given by Turchill, who was believed to be a clerk in the royal household, often said to be a clerk in the Exchequer. That is slightly confusing. The book which I use, rely on for this, is the account given in Poole's book on the Exchequer in the 12th Century. Poole was of course mainly interested in the Exchequer, but the Exchequer is rather misleading, because although an abacus of this kind was undoubtedly used in the Exchequer, and indeed gave its name to the Exchequer - the chess board - it was a simpler kind of abacus for a quite different purpose. It was for counting the sums of money that the sheriffs brought in from the counties, and it was only really...the sums that were done were only really additions and subtractions. But Turchill I think must have been a clerk in a more senior position, if you like, and he wrote about the abacus in general, and his account tallies very much with the one that was given in the manuscript described by Yeldham.

So, what were the two distinctive features? The grid was arranged in the columns, and the counters were labelled with Hindu-Arabic numerals. This is - Gerbert was around the turn of the millennium. In fact, I think he was Pope at the turn of the millennium. These accounts are 1100, beginning of the 1100s, the 12th Century, so this is all quite some time before even Fibonacci, and my conclusion is that, in the higher realms of government, and perhaps in the large monastic estates, Gerbert's abacus, with all the frills, was well-known in the 12th Century. Now, that shouldn't deflect us from noticing that of course older forms of the abacus, with plain counters, useful for doing additions and subtractions and accounting and so forth, were used by merchants and, as Fenny has said, we know that that carried on into I say the 17th Century - you may say the 18th Century. People still used the jetons and the counters for doing simple arithmetic. However, from a mathematical point of view, the real interest is that Gerbert's abacus evolved into a system where you weren't just moving the Hindu-Arabic numerals on the counters around, but actually using the same procedures, the same algorithms, for doing the calculations with the numerals themselves. In other words, what we would often call pen-reckoning, but doesn't have to be with a pen. It could be traced in the sand with a stick, or it could be they were written on a slate, and for both of those reasons, it's unlikely that we'll find many extant examples of them, but nevertheless, my belief is that the Gerbert

abacus and the pen reckoning algorithms, using Hindu-Arabic numerals, were very similar, and in certain quarters, they were well-known by this time.

Here is one of the few examples of pen-reckoning, from 1320. These are just long addition sums, but it's clear that that's using the sort of standard methods. There are the carrying symbols and so forth, and the numerals are almost recognisable to a modern eye there, so you can check that out.

So, one reason why I think there's some confusion about what was going on, about the technology that was available, is that there are misleading clues. One of the misleading clues is the oft-quoted fact that in 1299, the Florentine bankers' guild is said to have banned the use of the Hindu-Arabic numerals, and you will read this in a number of standard texts without really any further comment, but of course they weren't actually banning the use of the Hindu-Arabic numerals in the banks. All they were saying was that you must carry on publishing your accounts in the traditional method, using the Roman numerals to record these things, so that everybody, or at least everybody who was interested, could actually understand them. In other words, the general population who could understand Roman numerals could see what was going on. They were not banning the use of Hindu-Arabic numerals in the banks, and in fact, as I say there, almost certainly, for about 100 years at that time, the Hindu-Arabic numerals had been in use. And just to echo again something that Penny said, we know that this led to the establishment of the schools in Florence and other Italian city states, where the children were told - confusingly it says abaco, abaco and algorismo, but one shouldn't think of those as being distinct things. The algorithms for doing calculation is really what's being concerned there.

I recommend this book by Alexander Murray, which I don't see quoted very much in the history of maths literature. He approaches it from a medieval historian point of view, but in fact, he has several chapters which I find compelling and convincing on this subject.

Ah, there we go! You've seen this one before, and like you, I think it's a source of confusion. We have this smiling Boethius it's supposed to be, which of course is complete nonsense historically, doing his Hindu-Arabic sums there, algorithms; and we have poor old Pythagoras, looking very glum, trying to use an abacus, but he's using the plain counter sort of abacus and not enjoying it very much, by the look of it. So this is to be thought of as a sort of modern Bostock and Chandler sort of book, which sort of tells you how to do things at a certain level but doesn't actually lead you on to the higher levels of thought.

Okay, well let's see what happens when we try to put these two things together?

Let's now sort of see when we try to put these things together. Now, there's a problem here, that in order to describe this succinctly, to a modern audience, I'm more or less forced to use modern notations, symbols and so forth. We must not think that this was the way the medieval arithmeticians thought about it. They didn't even have an equals sign - as we know, that was Robert Recorde 200 years later. They certainly didn't have the idea of let x equal such-and-such, and then doing elementary algebra. So it was all done in a more...verbal kind of way, and the arithmetics of the time would talk about the rule of three, for example, which to a modern eye, is just a simple proportion sum, but they would have a number of mechanisms for working out, if it was a is to b as c is to d , depending on which one was the unknown that you wanted to calculate, you had different rules and methods for doing it. Okay, but I'm going to talk in these terms and hope that we can thereby elucidate what was going on, without overriding the mode of thinking that lay behind it.

So, by 1300, of course there were many places, centres of population, where trade was flourishing. Let's first of all think about a single place, x , and in this place - we'll go back to Jevens' distinction - there would be a money object, which I call M_x and there would be accounting units, which merchants would use to keep their books. The relation between the two would be determined by the authorities in that place. If it were London, even in earliest times, they were subject to the king for the coinage and so forth, so the king would say how much a money object that the king had issued ought to be accounted for, what it was worth, if you like, in accounting units.

Well, let's look at London. So the accounting unit was the sterling penny, which is just an abstract object from this point of view. It's an accounting unit. Now, there was something called a penny, a coin, a money object, but it wouldn't necessarily be the same as an accounting unit. It might have been clipped, and if you wanted to make sure you were getting your money's worth in any particular transaction, you might have to weigh out the silver coins and make sure you had the right weight of silver, and even then, you would be trusting that the coins were of the correct fineness, that is contained the correct amount of silver. So for larger transactions actually, in London at this time, things were calculated in marks, and a mark was said to be worth 160 sterling pence, or if you prefer to use the next accounting unit, 12 of these

sterling pence was an accounting unit called a shilling. There was no coin called a shilling - that's something that's quite clear - no coin called a shilling until the 16th Century in this country, but the accounting unit of a shilling goes back much further. So the mark is not a pound or a certain number of shillings; it's a certain number of pence - it's actually 13 shillings and 4 pence, a very odd amount - and when, in England, they started to mint gold coins, they tended to be actually in these proportions, nobels of 6 shillings and 8 pence, and things like that.

In Florence, on the other hand, they had of course a different kind of penny - denari piccoli - small pennies, and they were small, and they were often not made of silver - well, the objects were not. The accounting unit, however, which was the abstract thing, was - the accounting unit that they used was called the denari piccoli. By this time, we're talking about 1300. The gold florins were in use in Florence, and it was decreed that a gold florin was worth 348 of these accounting units - that's actually 29 twelves, 29 twelve times...yes, 29 times 12, yes, that's right. So the Florentine shilling was called a soldo, and it was 29 of those.

So suppose we're trying to send our shipload of wool from London to Florence. Obviously, the transaction is going to be governed by what we would call an exchange rate, and I'm going to tell you how I see this from a modern viewpoint, and then we'll look at what the Florentine bankers wrote down about this.

So I would define the exchange rate as the number of florins that equals one mark; in other words, it's the relationship between money objects, actually coin - well, actually there wasn't a coin called a mark. There might have been 160 pennies called a mark, but in fact, there might have been a weight of silver called a mark, or there might even be an ingot of silver called a mark. The use of ingots is quite well documented. So there's that exchange rate, which is the modern way of defining it, and of course there's the reciprocal one if you're going the number of marks that equals one florin and so forth. Now, this is not a constant number, fixed for all time, because it depends on economic factors, which is a sort of catchall term for saying that if there was a shortage of gold or a shortage of silver, gold and silver are used for other things rather than making money objects, and therefore they have a value which is affected by shortages and so forth. Other things as well can affect the exchange rate. So this is variable, and if we're going to do any sums involving the exchange rate, we've got to have a range of exchange rates available so that we can work out which one to use on a particular occasion.

This is what Pegolotti did. This is his table for the exchange between London and Florence, and it gives a range of values of a parameter alpha, and the corresponding values of a parameter beta.

So let's just look at one of the lines, focus on the one which has 34 in it, okay. What that says is that when 34 sterling pennies go for a florin, then the mark, the sterling mark, goes for 6 lire, 16 soldi, and 5-and-11-seventeenths denari. So each line is a statement of the form that I've written at the bottom. You'll notice there's the confusion...not confusion, but at least the double use. Some of the things here are the actual money objects that are going to be handed over, and some of them are the accounting units which are going to appear in the bankers' books. So that's why it has to be written in that form. So the exchange rate, e , that I had on the previous slide, has to be interpreted in terms of the Pegolotti table. So Pegolotti is making a table of values of a parameter beta in terms of parameter alpha.

Well, I think I've said one of these things already. The first one is that the number of denari in order to avoid having large numbers of denari, you have these higher units, super units, of the soldi and the lire, which incidentally, are the same multiples, 12 and 20, as would have been used in London, except that in London, it was a sterling penny, a shilling, and a pound. That's actually just a matter of convenience. We've also noted that the two types of money, or two uses of money, are signified in the table, because when payments were made, money objects had to be involved, and when the books were kept, the accounting units were involved.

Now, of course, we're already getting into the stage where we see new types of money being created, because the bill of exchange that I talked about is now becoming an object of money in itself, because it doesn't - if you'd got a bill of exchange which said somebody was going to pay you something, you didn't actually have to present it for payment yourself, you could try trading it on to somewhere else. So we already begin to see this layer upon layer of financial operations beginning to take place.

Here are the sums, and once again the caveat about this is not the way a Florentine banker would have done it, but it tells us in modern terms what the sums were, and there's some point in doing that. So, first of all, if the alpha in Pegolotti's table was the number of sterling pence that made a florin, and we can translate that into the number of marks that make a florin, and in terms of the e that I had before, that was the reciprocal one over e . Similarly, the beta, the answer that comes out - so the beta thing, to remind you, is the number of Florentine units, florins, that are equivalent to a mark. That

comes out to be e time vf , vf being the 384 number, the number of denari piccoli that make a florin, and eliminating, doing the algebra, which of course medieval arithmeticians didn't do it this way, you see the relationship between β and α straightaway. So in order to do this, make up his table, Pegolotti had done this sum, and that's why those strange fractions came into it. He had divided vl times vf , and vl - I've done this in general because the other thing perhaps to keep at the back of our minds is the fact that exchange was not just between two places. There was a whole network of places between which exchange could take place, and so we would make suitable substitutions for London and Florence, and we'd get the exchange between Florence and Bruges and places like that, for example. So, that's what the table is. It's a table basically of division sums.

But now, we can begin to see a little bit more what's going on - hopefully this has thrown some light on that. So, if you're a clerk in the bank in Florence, what you want to know is how you're going to account for the shipload of wool, which is supposed to be worth so-many marks, and what you're going to put in your books, and the answer to that... Well, no, first of all, the answer to that depends upon what is the exchange on [London], what is today's exchange rate, the parameter α . So Pegolotti's table told you that if α was 34, then β was whatever it was, okay, and in order to get the accounting answer that the clerk required, he would have to multiply the number of marks by the value of β , the value of β that corresponded to the given α . So the clerks were doing multiplication sums - relatively simple, probably could be done with the abacus, the unnumbered abacus, the abacus with plain counters. Certainly, the later books tell us how we could have done that kind of thing. But in order to draw up the table, either Pegolotti, or the person who calculated it for him, had to do a division sum, and quite a complicated division sum. He had to divide this number by α in order to obtain the correct value of β . So, this seems to me to make it very clear that even by this time, there were two different levels - perhaps more than two, but anyway, at least two different levels of expertise in arithmetic. You had the level of Bostock and Chandler calculus, where people can do the sums - in other words, when you know that the derivative of x squared is $2x$ or whatever it is - and you had some other people who really understood a bit more and who could do more complicated calculations. The second group of people, who would be, if you like, in the back room at the bank, they would be either, by this time, probably not using Gerbert's abacus, but actually using the Hindu-Arabic numerals pen-reckoning and so for, but their expertise, their methodology, had developed from the Gerbert version.

So, to round off... There are a large number of conclusions, and if you go to that book by Alexander Murray, which I mentioned, you will find a very useful attempt to put all this in the context of the growth of numerate thinking in the later Middle Ages, asking such questions as to what extent the growth in mathematics - Cardano, Tartalli and so forth, the solution of algebraic equations - to what extent that depended upon the base that had been established earlier on for commercial reasons. Of course we shouldn't think that commercial arithmetic was the only stimulus to the development of arithmetic. Astronomy was equally important, and the calculations that had to be done by the astronomers would have used, I believe, similar sophisticated methods and certainly weren't done with the plain counters that the merchants used.


So, well, let's just, as I say, summarise what's here. Different levels of expertise, and it's not always apparent from the evidence that's in front of us. The evidence can be unbalanced because we have the printed books from the 16th Century, which are telling us all about how merchants can do their calculations and so forth, and there's a lot of that. Evidence for the pen-reckoning method is harder to come by because it was essentially ephemeral, and whether it was done by scratching the figures in the sand or by writing them on slates, it tended to disappear.

And then finally, the generalisation about the numerate thinking, that there is in fact evidence that the commercial considerations were important in developing numerate thinking, possibly only indirectly, however, in that the abacus schools and so forth led to generations of Italians in particular, but other countries as well, where numerate thinking and so forth was part of the training. There were some things that didn't happen, and it's always...the dog in the night-time is always an interesting speculation. It doesn't seem that there was any serious progress in the concept of number at this stage. You might think that, given the fact that the exchange rate, for example, small changes in the exchange rate could lead to quite significant differences in the way the sums came out, you might think that that would lead to the idea of the number as a continuum, but in fact, the standard method throughout the Middle Ages of dealing with smaller quantities was to invent smaller units, so you had farthings, and then, at one point, they invented something called a mite, which was a fictitious thing equal to one twenty-fourth of a penny, but there was no idea of what would lead to the decimal notation, that is, having tenths, and then tenths of tenths, and tenths of tenths of tenths and so on, having the same multiplier over and over again. Although that was, in the Hindu-Arabic system, used for multiples, it didn't seem to come in for fractions, sub-multiples, until the 16th Century, as we know, and without that, it's hard to lead on to the idea of the number continuum and the infinite decimals and so forth which you require to deal with that, and of course the calculus itself, which requires the notion of small change to be codified in some way.

So, I hope that that's shed a little light on the topic. I'm, as I say, I was heartened by the fact that your talk had come to

similar conclusions in some of these aspects, and I'd be interested to hear any comments that people have. Thank you.

© Norman Biggs, 2008



25 APRIL 2008

WHEN COMPUTING MET FINANCE

DR DIETMAR MARINGER

Good afternoon. I think my talk differs a bit from all the other presentations in at least two respects: for one, in computational finance, we usually don't think in centuries, we think in terms of years and decades, and this already is quite a long time. One of my students came to see me the other day, and he asked me what do I think about this, his words, "ancient book", and he gave me a book from 1991! So we have some sort of different idea of what is a long span in computing, in particular computing and finance. Obviously, computing in itself has been around for quite some time, but in finance, it's a little bit tricky, because - this is probably the second thing which is different to some of the other topics - computing and finance, it's sort of a strange love affair, because it's one of the things where, initially, neither of them admits, yes, we do have common interests, and once you can no longer hide it, everyone says, oh obviously, what's your problem, it's always been joint interests! This is exactly what's happened in computational finance. So, for mathematical finance, for example, there are clear papers and clear dates where you can see, now, this is where Black Scholes came up with their idea, but in computational finance, sometimes it's difficult to pinpoint when something changed. One of the few areas where you actually can pinpoint something is when you look at institutional aspects.

So traditionally, stock markets worked in a rather market-type way, as you would expect a market. People gather, some of them want to buy, some of them want to sell, sometimes the roles change in between, and a traditional case for a stock market was something like a open outcry. So people, like in the picture, meet on a marketplace, or on a trading floor, and they just shout what they want to do, shout the prices and the market maker or they themselves find out what the prices are.

Some times, in the 1970s, '80s, markets and stock markets switched over to electronic markets, and this was something obvious, because when you look at different market places - the first one was in New York, the NASDAQ, which in 1971 opened its doors, and it was, from day one onwards, an electronic market. The London Stock Exchange closed its trading floor permanently in the early-1990s. There had been a parallel system for two or three years around, but in 1992, they went electronic. Strangely enough, Switzerland, revolution started slightly earlier. In the 1960s, there was the Swiss...the stock exchange [?]. In the 1980s already, they had very strong computer support, and in 1996, which is later than London, they also switched to an electronic market. New York followed only last year, where they currently have this hybrid market. So you can have both things, and you still have the trading bell which opens and closes the market obviously, and where still you have your people meeting in the room. So, when it comes to institutional aspects, there are some clear dates you can attach to certain things.

Another thing which probably made a difference to finance, as far as computing is concerned, is the advent of the internet, which an impact on two levels. For one, it provided information. So in the early 1990s, an internet browser looked something like this, where you had your very basic structure. You had your hyperlinks, and the nice thing was you, yourself, in particular if you work in the proper institution, could provide information which is open to everyone who has access to the internet.

So along came Bloomberg. This is a screenshot from 1996. Unfortunately, a couple of the pictures are missing, but

providing information was crucial in those days and had a real impact on trading behaviour.

Same for Reuters, and what I like about this screenshot, if you have a very close look - probably you can't read it, down here, it says "text version only". Those were the days where you really struggled with your broadband connection because broadband didn't exist as such, so you had very slow connections and you only got text information and you actually could choose whether you want to have pictures on it or not, or if you just want to have the text literally.

But the internet also has another, or had another impact: people started trading over the net. So it was not just professional, but it was also the average man in the street who could trade him or herself using the internet. So Ameritrade - again, this is a screenshot from 1996 - were amongst the first ones to provide these sort of services, and they pride themselves on this front page that they have over 300 million in assets. Nowadays, no one would be really impressed with this sort of number, but in those days, it was quite a big deal. Eventually, they also provided research tools. They registered their slogan "Believe in yourself." If you look at the date, this was late-1990s. After the burst of the internet bubble, this slogan was nowhere visible, but they still provided still this information, this cheap, relatively cheap access, so you could trade for \$8 per trade, provided you trade more than 10,000 stocks, but still, it was reasonably cheap in those days, because in those days, you had large margins. This is one of the aspects where computing really made a difference.

Just one last example, ICAP originated as a merger between companies in 1998, and now are London-based, next to Liverpool Street, and as far as I understand, are currently the largest internet brokers worldwide, and they're still based here in London.

So the technical revolution, and to some extent, the internet revolution, had an obvious impact on finance directly. Where it was less obvious that there was actually an impact, and when it really got started, is with all the other aspects. So one of the aspects currently being a big deal is automated trading.

Automated trading means you don't have a human trader giving a buy or sell order but you have a machine giving a buy or sell order. Now, a couple of years ago, it was, allegedly, roughly 10% of the volume traded on the London Stock Exchange based on orders by algorithms or by computers. Two years ago, it was 30%, last year it was 40%, and for 2008, they estimate 60% plus. So automated trading has become a big deal, and behind most of these systems stand more or less sophisticated trading algorithms. Some of them are more or less straightforward, some of them are less straightforward, but they do have a major impact on finance, on how stock prices behave, and on how markets behave obviously. So the idea is, with all this automated trading generated by machines and generated by computers, that these buy and sell orders are given by machines, and these machines follow certain algorithms, follow certain rules.

The reason why people used this sort of automated trading are multi-fold. What [was this thing is] arbitrage? Arbitrage means you can make money for nothing. We already had this example in a previous talk today. As you might gather from my accent, I'm not British, I'm Austrian. We have Euros, so if I come over, exchange my Euros to British Pounds, and immediately would exchange them back to Euros in a different country without any [temporal] delay, and I'm left with more Euros than I started off with, then this would be a case for arbitrage, and this obviously must not exist. There are very straightforward relationships, for example, between exchange rates and limits between exchange rates, give and take transaction costs, which must not be violated, and machines are very quick in spotting these inequilibrium. So machines can be used in automated trading to exploit arbitrage situations, but then, as a consequence, they have an impact on the price, they drive the price back into equilibrium, and the arbitrage opportunity vanishes.

The next thing where automated trading is used is risk management and hedging. Hedging means you want to reduce the risk in a portfolio or in a single asset, or in any sort of financial investment, because you want to limit it and you want to build, literally, a hedge around it, and this quite often is done with options. Now, we already had a very interesting talk about options and how option prices evolved and what the underlying assumptions are, and we already had a very detailed discussion of the Black Scholes equation.

Now, this Black Scholes equation is one model to price put option. The example in the morning was about call option, meaning I'm allowed to buy something, I have the right to buy something. The put option is the equivalent on the selling side, so you have the right to sell a certain underlying asset at a specific point in time for a pre-specified price - the strike price. Black Scholes came up with...or published this result in the early 1970s, and they assumed, or made a couple of quite realistic, or reasonable, assumptions that we have this geometric Brownian motion, as in process good enough to [describe] what the process of the underlying stock is, and just to keep things simple, we don't have a dividend until

maturity.

Now, some - actually, the usual thing for stocks to be at least one dividend a year, so a couple of years later, Black came along and suggested a slightly modified version of the Black Scholes equation, where he deals with the case that you have a European put and you have one dividend until maturity. What does it do? It simply corrects for it. He assumes the dividend payment is safe, so it just discounts it, he splits it off the stock price, it leaves the remainder as the new stock. Very clever idea...

But it's still a European put, meaning we are still only allowed to exercise at this one specific point in time. Now, if you have a put option, then if you look closely at the price and price behaviour, there's actually some point in time where you would be quite happy if you could sell the underlying right away and you don't have to wait until maturity. So for puts, unlike for calls, in puts, it is the case that sometimes you actually want to exercise prematurely. The only trouble is things suddenly become a little bit more complicated, and MacMillan eventually solved the problem and suggested this equation, to price an American put - again, assumption no dividend underneath.

In the previous talk, Schachermayer was quoted in one instance. Schachermayer in those days was working in Vienna. Vienna was an interesting place to live and to work in, and particularly in those days, and I was fortunate enough to work at the same department as Schachermayer with someone called Fischer. Fischer, in those days, also was working on option pricing, and he extended the model and introduced the case that you actually have one dividend until maturity, and this is the option price and these are the parameters that go on. So I think you get the idea: you can easily grow and grow and grow the complexity of this product, and with it, you can easily also grow the complexity to compute the result. If you have a closer look at this equation, you notice that, here, we already have a bi-variant normal distribution.

We also derived a different pricing model for credit risk - tricky to tell nowadays, but those were the days! - credit risk where we wanted to price guarantees on loans. The idea was, because, in those days, option pricing theory was **the** topic to look at, we used results from option pricing theory, so what we did is we had a model where we priced it as an option, on an option, on an option, on an option, and so on and so forth. For every point in time where you have to pay your interest, or when your loan is due, you introduce one additional option, because that is one point in time where something could happen. So if you have one interest payment and one point where you pay interest and pay back your loan, you have an option on an option. If you had two interest payments plus redemption, you had option on option on option. If you had...you get the idea! The problem was, for every additional option, you have got one additional dimension in your normal distribution.

Now, solving this problem, again, the equation got longer and longer. Solving numerically and really number-crunching problem like this meant if you have a four-dimensional normal distribution in those days, basically, you pushed a button. You had a five-dimensional one, you had time enough to get yourself a coffee; you have a six-dimensional one, you can wait over the weekend; you had a seven-dimensional normal distribution, it took substantially longer; eight-dimensional, we estimated roughly 10,000 years! Because the computational complexity explodes, and this is already one of the crucial things about computing: it's not good enough to have faster machines, because what's the good of a machine that is 10 times as fast? What's the difference between 10,000 years and 1,000 years? If you do, in particular now, this high frequency finance, it's simply not working, so eventually you have to come up with more sophisticated algorithms which circumvent the problem in itself, or eventually, you just draw a line and say, now, that's the limit of complexity we can deal with. So in actual fact, these sort of modelling approaches eventually came to a halt.

There were a couple of alternative option types. There were Bermudan options, because Bermudas are right in between Europe and America, and if Europe is one point in time that you can exercise, and America is any point in time that you can exercise, then obviously Bermuda is a good name for a type that you can...where you have a mixture. So you have some window in time where you can exercise. There are other exotic options, with all sort of fancy...things when you can exercise, how the exercise price is actually computed or predetermined or found out, where you have a situation if you hit a barrier once time to maturity, then it's good enough you don't have to hit it at the expiration day. Many alternatives - also now we know that CDOs and CDO squares, which were one of the ingredients for the credit crunch and the whole crisis recently, they all gave us quite a sort of a headache. Unfortunately, we can't use all the beautiful mathematics because we don't necessarily get to a closed form solution. And the next thing we also have to take in mind, following Black Scholes, in option pricing, quite often we make the assumption that we really do have this geometric Brownian motion, which ideally actually we should have. Unfortunately, stock markets do not behave accordingly.

Now, this is a distribution of the daily returns of the Dow Jones over a quarter of a century. Those of you who work in

statistics might recognise that this one lies similar to a normal distribution, which is one of the ingredients for the geometric Brownian motion, but it's not really a normal distribution because it's too slim. If you have a very close look, you'll find a couple of outliers, and these outliers should have happened with a probability of one in seven million years. In actual fact, we had a dozen of them over 25 years. So it happened with way too high probability and this is why computing now uses - when you actually - when it actually comes to solving these option pricing problems, Monte Carlo simulation is used. So the idea is you use simulations of the underlying stock paths, you find out what the options would be worth if this really is the outcome, you do this over and over and over and over again, and then eventually you get an idea of the distribution of the terminal price, for example, of the option, and then you get an idea of what this thing should be worth today, because there you have much more...much more sort of flexibility in designing the underlying - you can have as many dividends as you want. The problem is we never know how good we are with this sort of simulation, so it's always a good idea to? And people in mathematical finance are still looking very hard into option pricing theory and to writing models for this, which in computational finance obviously are always like gold dust, and they are the...the margins we would like to hit.

Nonetheless, this whole theory can be used for automated trading, and this was actually one of the first applications in automated trading, and if you remember the previous slide, one of the applications was, the first one was arbitrage, and the second one was hedging. Now, one of the main things with options is that their price is really driven by the price of the underlying. So if we have the right to buy to something, then obviously - the right for a specific price, then obviously this right is more valuable if the underlying is more valuable. So if the price of the underlying goes up, then this buying option increases in value. At the same time, the right to sell the underlying decreases in value. So the put has exactly the mirrored hockey stick we saw in the morning in the call pricing problem. The nice thing about the approach by Black Scholes was that they take this into account and their approach, what they say in their model, exactly the change in the put option given that the underlying changes, and this thing is called delta. That's the first derivative of the put price with respect to the underlying's price. This is actually a quite helpful thing because if you know that if your stock price drops by one pound, and your put goes up by, say, 50p, then what do you do? You buy two puts and one stock, and the price movements offset each other. That's the idea of hedging, as simple as that.

Obviously, if you look at the graphs, and the delta can, since it's the first derivative, it's a tangent on any of these lines - they just differ in terms of time to maturity - you get different deltas. But nonetheless, that's the way it works, and that's the underlying idea of this whole thing, of this no arbitrage condition, so actually the circle closes. Unfortunately, it does not always work as nicely as we would like to see it work.

This is 1987, and the main thing happened in mid-October, which really gave us all a headache and what you can see is, is this big jump, downward movement in the price, and then obviously one of the main assumptions in the Black Scholes model is violated, because we no longer have a continuous price process, we have a jump downwards, and even worse, thanks to the joint effort of mathematical finance and computational finance, these... downward movement was accelerated because, in those days, everyone believed in Black Scholes, all these hedging strategies had hard-wired the delta in their automated trading system. So if they wanted to insure against drops in the underlying prices, they [justified] buying or selling signals of the corresponding options. Eventually, the market runs out of liquidity, and eventually, it's just a vicious circle. Now, I'm not suggesting that this is the only ingredient for this...event. There were a couple of additional things going on because, at one point in time, trading was stopped, liquidity was an issue, but one of the ingredients really was the automated trading which was not done properly.

I think we had one question in the morning - what happens if the optimisation - if everyone uses the same optimisation technique? That's exactly the problem - was the problem in those days: they all had the same hedging strategy. So the good thing is, by now, we have learned our lessons, and now these sort of things should happen no more because we know what drove these sort of events, at least accelerated them, and I will come to that in a minute, how actually overcome the problem.

The other aspects in computational finance and automated trading are you want to have superior predictions, and...yep, automated trading is a self-fulfilling prophesy in itself.

So the next thing where we might be interested in is we want to have superior predictions. Again, this is one of the areas where finance, at some point in time, looked over to what computer science does, and one of the areas they looked into was artificial intelligence. Artificial intelligence has been around for a couple of decades by now. It's, again, difficult to officially mark the day in your calendar when its birthday is, but in the early 1900s, cybernetics was setting out - 1920s was one of the century's? Artificial intelligence itself, the term was coined in 1958. Now, this is sometimes quoted as the birth of artificial intelligence, but in those days, people had different ideas about what is an intelligent thing. So for quite some

time, having a computer programme that can play chess was the ultimate thing to achieve in computational intelligence. Having a system that can do mathematical logic was the ultimate thing to do in artificial intelligence.

So one of the fathers of artificial intelligence, John McCarthy, who actually coined this phrase, he was working in mathematical logic and how to use computer systems to do mathematical logic. Alan Turing, then working in Cambridge, was also interested in what actually is intelligence, so he came up with this, what's now called a Turing test, where you can - which, in those days, was a criterion on whether something is intelligent or not, and his suggestion was, if you "speak" (inverted commas) to this machine and you cannot tell does the answer come from the machine, because it's sitting behind a curtain, is the answer coming from a machine, is it coming from a person, and if you can't tell the difference, then it must be intelligent.

A couple of years later, Joseph Weizenbaum, then at the MIT if I'm not mistaken, wrote a nice little computer programme called ELIZA, which did exactly the same thing, and story has it that he showed this programme to his secretary and asked her to play around with it, and she actually did, and eventually, he came back, wanted to ask her how she was getting along, and she stopped him and said, "Oh, don't interrupt me, this is personal!" What the thing actually did is it had a couple of buzzwords, so if it didn't recognise any of the words, it just said, "Tell me more about it". This was obviously very much encouraging for a person to keep on typing. If it had certain words which resembled, or had in its list resembling something like "holiday", then it made statements, "Oh, that must be pleasant," or something like this. So very simple rules, and it actually passed the Turing test, at least when applied by this one person. So this idea of what is intelligence is - has been a problem ever since, and we still haven't found a clear definition of what is intelligence, because whenever you make, or come up with a definition, eventually you reach this hurdle, and then, oh no, it's not really intelligent - we must make it tougher, because it's just bits and bytes inside the machine.

But one of the crucial points in artificial intelligence was yet another PhD thesis, because we had Bachelier today, we had Markowitz today, all PhD thesisists. Marvin Minsky also wrote a PhD thesis in order to get his PhD, and he introduced something which is, by now, one of the standard methods, neural networks, and I'll talk about this in a minute. In those days, obviously, he hadn't...machine like this. He used 3,000 vacuum tubes to simulate a net of 40 neurons. Meanwhile, and similar to Bachelier, he also faced some sort of criticism, because in his examination, people - he was doing a PhD in Mathematics - people struggled to acknowledge that this actually is mathematics, what he is doing.

Nowadays, we know it actually can be regarded as some form of non-linear regression, so it is, in one way or another, it is mathematics, but nowadays, artificial intelligence has moved on. We now speak more of soft computing, because we have given up this idea that as long as it has mathematical logic inside and it's based on clear, well-defined rules, it is intelligent, as long as these rules are really clever. Now, we have something like soft computing.

We also have a new term which is called which is called computational intelligence, no longer artificial intelligence. I remember I went to a conference in Cardiff, a couple of years ago by now, which was one of the first conferences which actually had computational intelligence in its name, and no one actually knew what makes the difference between artificial intelligence and computational intelligence, so they had a competition, and every participant was asked to write on a piece of paper and submit it to a ballot, and the winning suggestion was it's Welsh for artificial intelligence!

It's really difficult to tell what's the difference. The main or the core of the definition is it uses computational methods. So it's not the cleverness of the idea but you have a computer computing something which looks and smells like an intelligent being, but it's no longer claimed that the thing itself is intelligent - it just mimics, it simulates intelligent behaviour. Again, this is not a 100% spot-on definition, so don't quote me on this, but this is the main idea.

Neural networks are probably the strongest bit of artificial intelligence that have made it into finance. Now, how do neural networks work? The idea is, or the story at least is, it mimics the brain cells. You have some input. If the input is strong enough, the cell triggers a signal itself. So, small input, obviously too small, no reaction. A slightly larger input, send it into the cell, the cell is activated, and it also sends a signal. The thing is, you can have several inputs into one neuron that just add it up, and again, if the sum is strong enough, it activates, but you also can have nets of neurons, so not just one neuron, but neurons which are interconnected. One signal sends its signal into many neurons, and every neuron receives its input from several input sources, and eventually, neurons are sending to neurons. That's the basic idea.

So if we simulate a net like this, then this is what we get.

If we had other inputs, other inputs, the outputs would be different.

Now, what you can do if you use this artificial net of networks is you can increase or decrease some of the inputs. How do you do that? You introduce weights for these links, and here, symbolised by a thick line, you multiply it with a factor of, say, 10. If it's a thin line, then you multiply it with a weight of, say, 0.5. So you increase, artificially increase or decrease the signals, and what you want to do is you want to get an output which is as close as possible to what the output actually should be. So what can you do? You can use this thing, for example, since it is some sort of regression, you can use this thing here to model stock prices. So what could you do? You input the past history. You input what the market currently does, and out comes a prediction for today's stock return.

How do you train a network? You use all data, you play around - well, hopefully not play around, but you find your weight such that it would have worked as good as possible in the past. Then you know you have a working network, at least on historic data, and then you apply it for a couple of days by feeding in new information and new inputs. You can also use it for other sort of things because you can, depending on how inside these neurons you have activation functions, step functions, signal function, all sort of different functions are possible, typically between 0 and one, or minus one and plus one. You can also use it binary decision making. You can also use it for probabilistic predictions. So these sort of methods became very popular in the 1980s, and in particular, in the 1990s, because by then, people had the computational CPU resources to actually train the networks and do all the data mining required to get sound results.

So, in the 1990s, you suddenly found literally hundreds and probably thousands of applications of neural networks to all different sorts of financial problems. They used it to predict bank failures, by using balance sheet information, or information about the customers. These sort of methods or models worked quite well. Others used it for exchange rate forecasting. Yet another networks were used to spot trading signals, so past data were fed into the network, and out came a buy or sell signal for stocks, for foreign exchanges, and so on and so forth. So this artificial intelligence side became very popular because, for some miraculous reason, it seemed to work. From a theoretical point of view, this should not have been able to make any money, because if we believe in a geometric Brownian motion, if we believe in a normal distribution, prices shouldn't have a memory, and this is one of the crucial assumptions in all the underlying theoretical models. But apparently, they do have some patterns, and nowadays, in one of the leading journals - probably **the** leading journal in finance is the Journal of Finance. They used neural networks for this in the Journal of Finance, but they also use it to detect - and there are not too many, but a couple at least, papers on technical trading rules, because, to some extent, it still is a mystery why it actually works, because we are back to the original question - now, if it works, why doesn't everybody use it, and why doesn't the effect in itself vanish? To some extent, the effect actually does vanish, so with neural networks, nowadays probably, you have a little bit of a hard time to really make money. So nowadays, you have to come up with something more sophisticated. But nonetheless, they are quite popular, and again, from a mathematical point of view, a statistical point of view, there's just one sort of non-linear regression and that's the way they actually can be treated.

Another thing which we in finance now use and that comes from artificial intelligence is evolutionary computation, which ticks more or less all the boxes to qualify for soft computing, because the idea with evolutionary computation is you don't pre-specify an awful lot of rules. You just set up a rather vague system and let it evolve of itself over generations. And what this does is it uses the principles of natural evolution, and one of the pioneering methods was the one suggested by John Holland, in the 1970s - if I'm not mistaken, yet another PhD thesis, or at least linked to it - where the idea was pretty similar to what we see in biology. We have two parents, they mate, produce offspring, and the offspring inherits part of one parent and part of the properties of the second parent, and there's also mutation going on. The main thing here is we don't have something like the DNA. We have something even simpler - we only have a binary code. If we have two parents, if parent one is 0011, and the next one is 1001, then what is done is you pick one random point, you cut the two genes into bits, and rearrange them. If you want to have mutation on top of it, you pick one of the genes and randomly change it - or not to randomly change it, because in a binary world, changing means go from a 0 to a 1 and from a 1 to a 0, so it's pretty simple actually. The good thing is, as simple as this might be, it works, because what you do is you start off with so-called population of these strings. You generate offspring, and you just check is the offspring better than one or two of the parents or one of the existing solutions. If so, the chances are it will replace; otherwise, it will not replace it.

This is actually one of the...or referring to one of the publications in this Journal of Finance. Blake LeBaron is one of the leading figures in this sort of application and also in artificial stock markets, yet another application of computational finance, where he provides a set of technical trading rules - moving average, hedge holder, you name it - and the binary string represents whether one market participant uses a certain rule or does not use it. So if the first rule is, for example, moving average, then this 0 indicates that this trader does not follow this rule, but it follows the second rule, and not the third and fourth, but the fifth rule, and so on.

So we have one trader who might look like this. We have another trader who has a different [chain], another one, and another one, and another one, and then they combine their rules, and then their performance is tested against their offspring's performance. If the offspring generates a higher profit, then chances are the original ones are eliminated and the new ones survive, or the other way round. The funny thing is, it works. Not really much guidance, but it works.

Another thing which is based on this idea is genetic programming, which is the next step, introduced - and here we're already coming to what, again, my student calls ancient - we're coming to the 1990s, and John [Cozer], mainly, suggested an approach which is called genetic programming, because his idea was, now, this shouldn't just work for bit strings, it actually should work also to generate computer code.

So if you represent an arbitrary equation or an arbitrary formula as a tree - the left one is the sine of x plus x divided by 4; the second one is $3 \times$ the sum of 2 and x . If we have this same idea, and just recombine it, we might get a new equation, and a new one, and a new one, and a new one, and if we also have mutations and we randomly substitute this plus sign with a minus sign, for example, then we'd probably get yet another rule. This idea of genetic programming also got very popular in finance, because what can you do? You can develop trading rules, and this brings us back to our previous idea of automated trading. So this is another highly important bit in...in computational finance, where people try and generate trading rules.

This is one example, which definitely is not a historic example, because it's current work of a PhD student of mine, but this is a good indication that this is what the industry is currently using and actually has been using for quite some time, but it's also a good example that what the finance industry actually does is not always quite visible. It's sort of a secretive love affair still, in this respect, because just a simple example...again, this is work with a PhD student of mine. One of the major investment and broker companies, Worldwide, they offer one quantitative position per year worldwide, and it was my PhD student who got the job because he's working on these sort of topics, but he had to sign that he's not talking about his work with them, and he was not allowed to use any of his results, any of his data he worked on during the summer, because they wanted to have the exclusive rights. So this is what, at the moment, makes it difficult to pinpoint what are the issues in computational finance, because we know what we do in academia, but we do not quite know what the industry actually does. We have a rough idea, but now we know neural networks, hot issue, genetic programming, hot issue, but again, we don't have many dates where we can say "It started in 2003, because this is the first paper," for example. So papers on this, for example, have been around for 5 to 10 years by now, but many of the applications for GPs are in different areas, they are not in finance, but it is pretty obvious that many people in finance use them.

Another area where computing and finance met is optimisation. We had this brilliant talk in the morning about how optimisation changed the face of finance because, let's face it, without quadratic programming, Markowitz's problem could not have been solved. It required the idea of quadratic programming. So if you have Markowitz's problem, you can treat it with a quadratic programming approach. If you give up the idea of Markowitz that short selling is not allowed, and introduce short selling so that you actually have a negative code and a negative sheet, then it actually becomes [an alternative to a closed form] solution. The problem is the world is not always normally distributed. So it's not something like this, or if this - if I can show this just for a second...

This comes from real assets. We are again in the volatility and returns space, and we get this hyperbola or parabola, depending on whether you have variance or standard deviation as your risk measure. However, if you're looking at what do these portfolios do in terms of skewness and introduce this as your third dimension, then things suddenly become very messy, because suddenly it's no longer clear what you actually want to do, and what is really good, because suddenly, you have these outliers and you do not know how much of a positive outlier offsets me for many, many small losses, for example. So it's very tricky to come up with a good utility function. One of the beauties, real beauties, about Markowitz is you don't have to make any assumptions about your investors apart that they are rational, but you don't have to assume that you are very risk-averse, or not risk-averse. You get the basic result - this curve, regardless of the risk aversion, because this is something for low risk aversion person, this one is for high risk aversion person, but you don't need to know it when you optimise it. If you look at an element like this, you need to know what to pick, and the same is true in particular then if you start changing your weights. Then, suddenly, the thing might look completely different.

Again, the only thing you can do is you play around with the weights of your assets, and then, suddenly, you have no longer functions, because this thing turns into a curly-wurly, and this is not what you want to see in optimisation.

Another thing that happens is new risk measures have come along. Value-at-Risk, for example: Value-at-Risk is not the standard deviation of what you expect, but it is the lower quantile. This is actually quite close to what is the everyday

notion of risk, because the everyday notion of risk is that things go wrong, not by how much I deviate from my expected value, and this is what standard deviation measures - both upside and downside risk. If you use a normal distribution, it looks like this. If you use an empirical distribution, it looks like this. Now this, the thing is it's slightly difficult to optimise.

How can you solve it? You use, again, evolutionary methods, or other methods inspired by nature. Simulated annealing is one of these methods, where you mimic how, when liquids solidify, how crystals emerge, because they want, the particles want to arrange themselves so that energy is minimised, required to keep this state stable.

A pain in your kitchen, but actually quite clever when it comes to find shortest routes, the travelling salesman problem was mentioned, and lay pheromone trails, based on a reinforcement principle, and very quickly find the shortest routes between their nest and your sugar and candy box in the living room, so they are quite efficient at this, and we can use this for optimisation.

Another method is differential evolution. So if we look into this problem again, this is the problem we just had on the slide. Obviously, a traditional gradient-based search wouldn't get us anywhere, because gradient is like you drop a ball and gravity directs it down, but it very quickly will get stuck in a local optimum. What we used nowadays or what people use nowadays, evolutionary methods, where, again, these principles from evolution are used, where current solutions are combined and recombined and not so good ones are eliminated at an early stage. If you have a very close look, you can - we want to minimise our risk, then it's probably not a good idea to be in these high risk regions here. I'm not too sure whether this is not visible at the time, but in this case, the deeper, or the further down, the better it is, and if you have a very close look at the graph, then you recognise that this evolution drags the solutions very quickly to the purple areas and very quickly away from the high areas. So again, not much intelligence, actual intelligence, because it's not a clever [?], it's computational intelligence. It looks as if they move in the right direction because they know what to do.

So, I think I need or I should finish eventually. The one thing I definitely haven't achieved is to answer the question when did computing and finance meet, but probably I managed to shed a little bit of light onto the question of where they met. So they met in institutional aspects, they met in terms of pricing, they met in terms of financial management, automated trading, and - I didn't address this - they also met in terms of simulators in artificial stock markets, which you can use for policy design. So again, you build your little world, which behaves, hopefully, close to the real world.

What sort of methods have made their way from computing into finance? It's basically - the first thing was, obviously, hardware, and hardware-related things - information systems, databases, actual electronic trading systems. The next thing are efficient methods, so very similar to the presentation this morning, having efficient methods that can solve quadratic optimisation problems or complex optimisation problems - not necessarily quadratic ones - are extremely helpful in finance and are widely used in terms of toolboxes or tailor-made software. Optimisation is a hot issue, but also artificial intelligence is a hot issue. But, the further down we go, on this line, the more difficult it is to say what actually is going on in the industry. It's a little bit easier to say what's going on in the literature, so if you have a look at the literature, you get an idea that this really is what people do, but once again, it is sort of a secretive love affair.

© Dr Dietmar Maringer, 2008