Cancer can give you Maths!

Professor Philip Maini FRS MeedSci

1 May 2008

I would like to start off by thanking Gresham College for this invitation. I am from the Mathematical Institute in Oxford. This is a place called the Centre for Mathematical Biology, which is about 25 years old, and the Oxford Centre for Integrative Systems Biology in the Biochemistry Department, which is one of a number of centres that have been set up recently, as is a part of a huge growth in this area.

We work on many different aspects of the applications of mathematics to biology. For example, developmental biology, how structures form in the embryo, wound healing and scar tissue, but what I am going to talk about today is the modelling of cancer. I am not going to talk about any particular cancer, because the work that we do is to try and understand fundamental ideas associated with the disease, how fundamental processes interact with each other, rather than targeting one particular cancer.

To give a very brief overview, the body is composed of cells, and these cells are dividing and do various things, so there have to be tight controls on the system so that the body remains in a steady and homeostatic state. So for example, cells have to be able to respond to various signals that tell them to stop dividing, because maybe they are dividing too much and if you keep on dividing, then a problem will arise. Also, cells are instructed every so often to kill themselves, in a process called apoptosis - cell suicide. So obviously, if a cell somehow does not respond to these two cues, for example, it continues to proliferate even though it is receiving signals that are telling it to stop, or if it stops listening to the cue that says we have too many cells now, you have to kill yourself, then obviously what is going to happen is an uncontrollable growth. So when I said that the cell stops listening or responding to these signals, what I mean is that mutations occur inside the cell, and so the cell does not respond appropriately to these signals, and then you get a growing mass of cells that eventually leads to a tumour.

So one of the challenges in cancer biology, and indeed in any area in biology, is trying to link the different levels of scale within the organism. I have been talking about the cell level, where you can get a growing mass of cells which can cause a problem, but that problem arose due to something happening at the gene level, so something was happening at the genetic level that was changing the properties of the cell, and that was manifested at the very top level of the cells and tissues. So how can we possibly link going from the bottom to the top? The short answer to that question is I do not know, and nobody knows. Also, what makes matters worse is that the relationship is not just bottom-up from the genes up to the cells, but top-down from the cells down to the genes. This is because the genes lead to protein production which affects how the cells behave, and the cells, in turn, modify their environment which affects how the genes are regulated. Because of these factors, and because of the vast differences between cells, there is actually no mathematical theory that can deal with this problem. It has not been invented yet, and there is an open question as to whether or not one can look at this mathematically.

But what I am going to do is present some simple ideas, some of which are looking from the cell level upwards, and then, towards the end of the lecture, I will give a very simpleminded idea of hop to try to link

across the scales for some special cases.

So, if something goes wrong in a cell and you have mutations, then it starts to grow out of control. If you do this in the laboratory, you find that the cell mass will grow in a sphere, and when it roughly reaches one millimetre in diameter, you get a black necrotic core in the centre. Necrotic means dead cells, so necrosis and apoptosis are two different things, giving the same results, namely a dead cell: apoptosis is suicide, the cell kills itself; necrosis is when the cell dies because of environmental conditions.

The necrotic core develops when the mass reaches approximately one millimetre in diameter, and the reason for that is due to oxygen. The mass of cells is getting nutrient, and nutrient is oxygen, and the oxygen is diffusing, and the diffusion length scale of oxygen is one millimetre. In other words, once the growth get to roughly the size of one millimetre, it is taking the oxygen too long to go from the edge through to the centre, so the cells at the centre start to die.

There is a simple equation which comes from some rather tricky mathematics to describes this. If T is for time, L is length scale and D is diffusion, this diffusion equation looks like this:

$$T = \frac{L^2}{D}$$

From this you can conclude that a typical timescale at which something diffuses is linked to the length scale over which it diffuses by this formula, $T = L^2/D$.

So if we take, say, something like an amoeba, the length scale is, roughly speaking, 10-3 centimetres, oxygen diffusion coefficient is 10-7centimetres squared per second, and that means that the timescale of diffusion across an amoeba is, roughly speaking, ten seconds. That means if an amoeba wants to live via just diffusion, it can do so perfectly easily.

Take an elephant, suppose an elephant decided it was going to live solely by diffusing of oxygen. If you do the same calculation, it boils down to 3,000 years. So that means that if this elephant, a thousand years ago, decided that it wanted to have a meal of oxygen, it would get it only about now. Therefore, obviously in the meantime it would have died. This is why elephants do not use diffusion alone, and this is why organisms when they get to a certain size can only survive if they have a circulating mechanism to circulate nutrients around.

It is the same thing with the cancer. If you think of the cancer as an organism, when it is still very small, it can survive just with using oxygen diffusion, but when it gets larger, it can no longer survive by diffusion alone and so it has to do something else.

Therefore, this necrotic region in the centre of the cell is a hypoxic region, meaning that there is a low oxygen level. Because of this, cells around the centre region start to secrete chemicals, which are called tumour angiogenesis factors. These tumour angiogenesis factors diffuse out of the tumour, and if this is in a body, they will diffuse into the blood vessels. Once they are here, they cause the basement membrane of the blood vessels to break up, releasing the cells that form the blood vessels. These cells are chemotactic to the tumour angiogenesis factor. This means that they move up the gradient of tumour angiogenesis factor, so they start to move towards the tumour, because the tumour is where the gradient is highest. But the very aim of the existence of these endothelial cells is to curl up and form tubes, so they will form a vasculature network as they move towards the tumour, and therefore, you will get more blood vessels within the tumour, and this is the process of angiogenesis. This is a natural process that occurs in the body. Any time you get a deep cut that goes through into the dermis and affects the blood vessels, it

will heals up due to this process. In order to heal, more blood vessels have to come into the cut and have re-established the blood supply, and that is done via this process of angiogenesis. So now that you have a situation whereby the tumour has now got all these blood vessels within it, it can start to grow. It has got lots of nutrients, and now it can start to grow. I should say, every single thing I have talked about has been modelled mathematically, but I am not going into the details of the modelling.

At this point the tumour encounters another problem, because it is not growing into open space but has normal cells there. So it has got to compete with the normal cells for nutrients and for space. Here, there is a paradox, because the acidic environment of the normal body favours the normal cells, so the normal cells are better adapted to the environment of the body than the tumour cells and therefore should win in this competition. Moreover, sometimes it is noticed that the tumour cells undergo a form of metabolism called anaerobic metabolism, which is twenty times less efficient that aerobic metabolism. So, you have got an environment that favours normal cells and the tumour cells are producing energy that is twenty times less efficient than these normal cells, and so in that competition it seems obvious who is going to win. But, as we all know, sometimes does not happen and the tumour cells win. So how can that possibly happen?

Well, one of the by-products of aerobic respiration is lactic acid. We have all experienced the effects of lactic acid, because any time you have gone for a run and get cramp in the muscles in your leg, that is because your muscles are running out of oxygen, so they undergo a different metabolism, which is the anaerobic respiration, which produces lactic acid. Then of course, you stop running and you rub your leg, and after a few minutes, the lactic acid disappears and everything's fine. But with cancerous tumours, if these cells are constantly producing lactic acid, that actually kills the normal cells, because it is toxic to the normal cells, and this creates the space for the cells to move in.

So about 15 years ago, Bob Gatenby, a radiologist, was working with a physicist, Ed Gawlinski, and they came up with a model to describe this. I really like this model because it was first proposed almost 100 years ago to describe how you have populations competing with each other. It is the Lotka-Volterra model for how predators and prey interact with each other. If we take N1 to be the normal cells, N2 as the tumour cells, and Las lactic acid, then we can give the model in the following formula:

I am not going to go into the details of this, but what these equations are staying is that the tumour cells, N2, produce lactic acid. Lactic acid diffuses into the environment, where it kills the normal cells. That then creates a space, and allows these tumour cells to diffuse into the space.

These types of equations have been studied before and it is well known the behaviour. It features in overlapping populations where you get the good competitor moving in and pushing back the bad or weaker competitor. Because this diffusion coefficient is a bit different to the usual diffusion coefficient, what they found was that the model predicted you had a gap between the advancing front of the tumour cells and the receding front of the normal cells. So Gatenby, who is a radiologist, went to do some experiments to see if this gap really existed, and he found it. Indeed, when you look at samples of cancer growths in the body, you can see that there really is this gap between the normal cells and the tumour cells. I think this is a very nice example of a very simple mathematical model giving a prediction which is then tested and shown to be true. Obviously this suggests that what one should be doing is trying to manipulate acidic levels to try and control the tumour.

One of the things that Gatenby found when he did this model was that it gave two types of behaviour: either the tumours always invade; or they died out. So he did not get a situation of a benign tumour, a tumour that grows a little bit and then just sits there. This is where we mathematicians got involved.

Some of my group got involved in this and wrote down a very simple equation to describe this process,

which you can solve very easily. What you find in doing this is that as the tumour grows, the tumour is producing lots of acid, which is also toxic to the tumour, so when the tumour reaches about one millimetre in diameter, it starts to poison itself. You will remember that one millimetre in diameter is roughly the size a tumour needs to grow before it runs out of oxygen. So now you have got the situation whereby the tumour is running out of oxygen and it is poisoning itself, so it needs a pumping mechanism for two reasons: to pump food in and pump the poisons out. So this is an alternative explanation for why it needs angiogenesis.

If you put the growth into the system, then you see that the tumour will grow exponentially. Gatenby could get such results as these, but he could not get a result which saw the tumour kept at a steady state, which the mathematics predicted as possible.

If you are a mathematician then this will suggest to you that there is a bifurcation here. This means that, as I vary one parameter, the qualitative behaviour of the system changes. Therefore, we want to vary that parameter so that we get the result we are looking for. We thus go back and look at that parameter to try and figure out what biochemistry is encompassed in that parameter. We then try to vary it in the way predicted by the model to get me to the desired result. This is precisely what our experimental collaborators are doing at the moment, to see if they see this behaviour.

When we put the model into a normal environment, it predicted a massive gap, but the gap that we observed in our experiments was not that big. The reason for that is that in this model we have included proliferating tumours cells which are dividing, and dead cells, but in fact, what you really have is you have a rim of proliferating cells and then you have a region of cells that are quiescent. The cells at the inside the rim are not dying and they are not dividing - they are dormant. So the nutrient concentrations and the acid concentrations are such that it is not good for them to divide, nor is it sufficiently bad for them to die, so they are just quiescent.

Because these quiescent cells do not produce acid, we are overestimating the amount of acid that is being produced because we are effectively saying that there is a big layer of proliferating cells and then a core of dead cells. Really, there is a thin layer of proliferating cells, then a layer of quiescent cells, and then dead cells. If you put that into the model, what you find is you get a gap that is precisely the size that was shown by the biological experiments.

With this finding we are beginning to see that acid seems to play an important effect in the invasion dynamics of a tumour. This enables us to think that we really ought to explore a bit further and it encourages us to look into the effects of acid.

So, to go back to biology, we can consider a duct in the breast; a ductile carcinoma in situ. Normally what happens here is that you have got a basement membrane which cells just line. But when there is a tumour present, what you see happening is that the tumour will grow inside the duct and then will invade out of the basement membrane and invade the tissue. That is where you get metastasis and those are the ones that are fatal.

So how does this all come about? Normal cells need to be attached to the basement membrane in order for them to survive, so if they become detached from the membrane, they will die. So if these cells are going to invade into the duct, something must happen to the cells that enables them to move away from the membrane and still survive. This change away from what is normal is a mutation.

Suppose that happens to the cells and these cells start to grow, but then they run against the problem of not having enough oxygen after a while, because they can only grow in a certain amount. As we have

mentioned, oxygen can only diffuse a certain distance, so the cells will begin to die, unless they start to use anaerobic respiration, which does not require oxygen. Therefore, they need to undergo some sort of mutation that enables them to use anaerobic respiration. But if these cells develop the possibility of growth without oxygen, we need to remember that the by-product of anaerobic respiration is acid. Because of this, the cells will start to kill themselves unless there is another mutation that enables the cell to survive.

We did a mathematical model of that and produced a hybrid cellular automaton model. Cellular automaton is basically a matter of representing each cell by a unit, and you give that unit rules. These rules determine when it should divide, when it should die, when it should move, etc. We can say to the cell, okay, measure your oxygen concentration, and if it is very high, divide; if it is very low, die; if it is in between, just sit there and do nothing. So then we need to know what the oxygen distribution is, and the oxygen distribution is written down as a reaction to fusion equation, a bit like I mentioned right at the beginning. Then we can bring in acid into the system and glucose, and then we can say that every time a cell divides, there is a certain probability it will mutate. If this mutation makes it an acid-resistant cell, then it will survive it high acid concentrations, whereas its neighbour which is not acid-resistant will die when the acid concentration gets high.

You can put all these verbal ideas into the model mathematically, and then you run the model. This whole process is called somatic evolution. It is not a new idea - you will all have heard of evolution from the point of view of populations. Here, we are talking about a population of cells, so we are talking about evolution occurring inside the body, which that is where 'somatic' comes from.

So we did our simulations, and we had to do them many times as they are stochastic in the sense that there is a random probability of cells mutating, so you get different simulations every time you run them. But through these simulations as got the findings that, first of all there are the normal cells, then, at some point, a new cell type takes over and they are the ones they do not need to be stuck to the basement membrane to survive, so they start to move in. They end up being replaced by the glycolytic cells, because they cannot go in very far because they run out of oxygen. Then the glycolytic cells start poisoning themselves, and you get the acid-resistant cells.

We mathematicians are obsessed with something known as asymptotically stable steady states, so we always tend to look at the final conclusion of the thing. But, if you are looking at the asymptotically stable steady state, it is a pretty silly thing because the asymptotically stable steady state is death. But what is interesting is how you approach death.

When we looked at these results we were very proud of ourselves for it, but Bob Gatenby saw the results and was uninterested in these, to him, obvious results. What interested him was something we had not thought was interesting at all - the transient behaviour. This suggested to him that a growing tumour should show a mixture of phenotypes, with up-regulated glycolytic cells surrounded by normal cells and hyperplastic cells. These tests are on-going as at present, but the results are already showing what the model predicted. So what we want to do is to look into the detail of that to see whether there are quantitative predictions that we can create.

So far we have either had a situation where the cells are in a bath of nutrient, so that was the growing spheroid, or we had the situation we have just talked about, where the blood vessels were outside the duct and the cells were growing inwards. In both those examples, there was no nutrient flow actually within the tumour - nutrient was coming from outside.

What we are looking at in this system now is a situation where there is blood actually flowing into the tumour. For many years now people have proposed anti-angiogenesis treatments which target the new blood vessels and kill them, so that hopefully then the tumour cannot grow and it will effectively starve to

death. There was a huge amount of fuss made about this many years ago, and it has been very successful in mice. If a mouse gets a tumour, there is no problem at all and lots of biologists can cure the mouse. But anti-angiogenesis treatments have not been as successful in humans as they have been in mouse, and nobody understands why that is the case.

So people have looked at combination treatments. For instance, you can have anti-angiogenesis and chemotherapy. So if you want to do a model that looks at chemotherapy plus anti-angiogenesis treatment, you need to have a model that, at one level, looks at the tissue level, where nutrient is coming in, but, at another level, looks at what is going on inside the cell, because the cell cycle is due to various proteins that cycle.

So now it looks like you need a truly multi-scale model. But in fact, you do not need a multi-scale model because we already know the answer to this, because chemotherapy or radiotherapy depend on oxygen. Radiotherapy uses oxygen to kill cells. Certain chemotherapy relies on killing cells that are dividing, and a cell needs oxygen to divide. So in other words, these two treatments work really well if there is oxygen around. On the other hand, anti-angiogenesis works by stopping the blood supply to the tumour, depriving it of oxygen, and so the cells die.

Therefore, these two treaments should not work together, because through anti-angiogenesis I deny the tumour oxygen, but radiotherapy and chemotherapy rely on oxygen to do its job. So anti-angiogenesis treatment coupled with chemotherapy should not work, or, at least, it should not be any better than doing one or the other. But studies show that it is much more successful, so we know that something has gone wrong in our logical argument.

Because this just does not seem to make sense, a mathematical model would be of great help of it could help us to understand this process. Like a lot of things in mathematics, you go through loads of hard mathematics, and then when you understand it, it is pretty obvious.

In the working of this we need to create a multi-scale model, linking from very low scale right up to very high scale. You have got the vascular layer, which are blood vessels. Blood vessels are not like the pipes of plumbing in our houses, they are very dynamic. They change their diameter because they respond to cues determining whether more or less blood is needed elsewhere in the body. Also, if there is not enough flow of blood into a blood vessel, the blood vessel will disintegrate. This leads to a question of how this system should be organised. This is because, given that the oxygen diffuses roughly twenty cell lengths from a blood vessel, we need a blood vessel every twenty cell lengths. So how would you design such a system?

One thing you could do is make very careful measurements and then you could very carefully design it, or another way you could do it is to just slap down randomly lots of blood vessels, and you can just let it sort itself out. Because the blood is coming in, if it flows down one vessel more than the other vessel, eventually that vessel will atrophy. In early development, what you see is a lot of blood vessels to begin with, and then vessel pruning by the self-organisation, so the system adapts to the environment it is in, and that is a very robust way to solve the problem.

So what is happening here is that these cells are producing various cues that are saying to the blood vessels dilate so that more nutrient can pass through it. Then these cells are all dividing, they are undergoing mutations that are making them turn from normal cells to cancer cells, and when they undergo mutations, they then start to behave in a different way.

We can now go back to the model I just described, where we had our little cells as little units, and instead

of just giving the cell little rules, we have little models going on inside the cell, models that are tracking the cell cycle proteins and therefore telling the cell when it can divide, models that are tracking the production of VEGF and when this VEGF is an angiogenesis factor. Then, at the oxygen level, instead of just having the oxygen diffusing from outside into the system, we have got blood vessels coming in to bringing the oxygen, and then the blood vessels are responding to metabolic and mechanical cues, changing their widths, their radii. It is in this way that you are getting an interplay between things happening at the tissue level right down to things happening within the cell.

There are a lot of problems with this model. One of the problems is a problem actually with mathematical modelling as a whole. It is the question of where you get your parameters and your data from? Most data you get is from the healthy state, but we are not interested in the healthy state. We are not interested in the healthy heart, or the healthy tissue. We are interested in the diseased heart, the diseased tissue, and there is very little data available for such systems. Obviously, if you see someone has got cancer, you are not going to try and cure the person, not measure the development of the cancer. So many of the constituents that we use in this model are from normal situation, so there is the big question of whether this applies to the tumour?

It is known that the tumour cells respond differently to vascular cues than normal cells, but it is not known what that difference is. So what we have done is try to play around all three parameters and see what happens. This is another thing that is a big fault with this model. I am telling you all the faults here, but it's better to be honest rather than to deceive.

One thing this model gives us is the total amount of tumour as a function of time, and you can see it oscillates, which is what is observed on tumour loads. It is very easy to understand how that oscillation occurs. The tumour starts to grow, produces lots more cells and eats up all the oxygen. One thing I have not mentioned which you also include in this model is that it puts pressure on the blood vessels and causes them to collapse, so no oxygen goes in. Therefore, if no oxygen goes in, the cells start to die, they also start to produce VEGF. The combination of less cells and the production of VEGF, means oxygen levels go up, so cells can start dividing again and the number of cells goes up. This gives you this cycle.

Our mathematical treatment of anti-angiogenesis gives an interesting result. You will remember that anti-angiogenesis is a process by which you reduce the blood vasculature into the system and that reduces the number of tumour cells because it reduces the amount of oxygen that goes into the system. Through our model we found that the higher the density of blood vessels, the lower the number of tumours. So this model predicts that certain anti-angiogenesis treatments will increase the tumour load, not decrease it. The reason for that is that it is actually making the delivery of the nutrient more efficient. It is doing vessel pruning. So now we have a possible explanation for why anti-angiogenesis plus chemotherapy works; because anti-angiogenesis is giving you a window of opportunity in which the oxygen delivery to the system becomes really good, so all the cells start to divide, and now you wallop them with chemotherapy, and that is better than doing one or the other. Then, of course, if you continue to do chemotherapy or continue to do anti-angiogenesis treatment, you do bring the thing down.

Now you see why my logical argument may not be correct right at the beginning. This is because my logical argument made the assumption that efficient flow of nutrient was an increasing function of density of vessels. This was intrinsic in my verbal argument as it was what underpinned it, and yet it's proven completely otherwise in the mathematical modelling. This is why mathematics is so good compared to verbal reasoning. This is because, when you do verbal reasoning, you do not have to make your assumptions explicit, but when you do a mathematical model, you have to make your assumptions explicit.

As I am speaking to a London audience, I presume that you already know that flow does not increase with density, and the London traffic will attest to this. To help the flow of traffic through this city it is not as simple as putting in lots more roads. It might help, but it might end up that you go round in circles and not

get out the other end. It might be that there is an optimal density of roads that allows the flow to be perfect. Therefore, if you have got a system where the flow is not perfect, what do you do? Do you decide to decrease the number of roads or increase the number of roads? You need to know where you are with regards that optimal. If you are beyond the optimal you should decrease the density to increase the flow. This is good for roads, and not good for tumour. But if you are below the optimum point, you should decrease the density and the flow, and then you kill off the tumour.

So what we are doing at the minute is working with some experimentalists that are actually trying to test whether this bi-phasic response of flow against density actually exists, or whether it is just an artefact.

So, to summarise what I have talked about today: I started off talking about the effects of acid, and the effect of how that might mediate invasion of the normal tissue by tumour cells by the tumour cells. Then I extended that a little bit further to include mutations and I talked about somatic evolution, and then looked a bit more at the effects. We have tried to set up a model architecture that allows us to explore the interplay between processes that are acting on very different length scales. What you realise from this is this modelling is very simple and very crude, and wrong! Perhaps all models are ultimately wrong, but some models are useful. What we trying to do with this model is to use it as a hypothesis-generating mechanism. As I mentioned, the models that we have done on the acid actually caused an experimentalist to go and do various experiments to test various hypotheses.

So I hope I have given you a sense of the sort of mathematical modelling that we are looking at in terms of tumours. There is also a huge area of modelling concerned at the genetic model, and although I have not mentioned any of that, that is very important stuff, but I am not involved in that area. Our expertise is in this type of modelling, of dynamics and looking at the consequences of mutations and genetic behaviour, rather than the genetic behaviour itself. But that is another huge area for research which is currently going on, but I must be satisfied with what has already been covered in this lecture.