Mathematics and Climate Change

What can a mathematician do about climate change, and why should one do something?

What insights do we bring?

What can attract us to work in the area or to learn about it?

Is it any use?

Can we combine our passion for truth and abstraction with our political beliefs?

Can mathematics go together with politics?

I don't mean politics in the sense of party politics, in which one advocates a particular line. I don't think that is what politics should be about. In English we have the two words *politics* and *citizenship*, which come, respectively, from the Greek and Latin words for city. Citizenship is generally regarded as worthy but dull; politics, especially at the moment, is exciting but disreputable, with a strong tendency to immorality.

My trying to convince you of my *political* views would be regarded as illicit, an abuse of my position; trying to recruit you to a particular view of citizenship would more likely be regarded as dull and quixotic, and perhaps slightly eccentric. Nevertheless, tonight's topic cannot avoid at least paddling in these dangerous waters.

The key difference between citizenship and politics is that citizenship implies cooperation, while politics centres on conflict, rivalry, authority and power.

Both clusters of ideas are susceptible to a mathematical treatment, through *Game Theory*.

Game Theory was invented by a mathematician, John von Neumann, and an economist, Oscar Morgenstern, (whose mother was the illegitimate child of Frederick III, the Georman emperor). von Neumann was also responsible for the introduction of the formalism of operators on Hilbert Spaces in Quantum Mechanics, and a significant contributor to the development of ther hydrogen bomb. Their 1944 book, *Theory of Games and Economic Behaviour*, is generally regarded as the starting point of Game Theory. For a modern and very short introduction, see Ken Binmore's book [2]. What is Game Theory? Roughly speaking, it is the description of the strategic considerations intervening in the interactions of individuals. Typically, in the somewhat abstract and simplified scenarios it envisages, individuals are faced with a number of choices, each of which leads to various consequences. The game is governed by a "utility function" or "payoff": each combination of choices determines a payoff for each player, and under the assumption of some degree of rationality, each player attempts to maximise his payoff.

Let me give some examples. There are in fact a number of classic games which appear at the start of every introduction to Game Theory. The first of these is called the *Prisoner's Dilemma*. It is very simple. Two suspects are arrested. They are in fact guilty of a crime, but in order to raise the tone of the discussion let's make them political prisoners instead of common criminals. They are interrogated separately; neither has any idea of what the other will say or has said. Each can either defect (confess), in the process implicating the other, or keep silent. The payoff in this game is a negative one: years in prison.

	Defect	Silent
Defect	-9 -9	-10 0
Silent	0 -10	-1 -1

The biggest payoff is if I defect, but my foolish partner in crime remains silent. For turning state's witness, I am released without any punishment, whereas my partner gets 10 years in the slammer. If he also defects, we each get 9 years. If neither I nor he defects, we are convicted of a lesser charge – downloading a book on Game Theory – and sent to prison for one year each. What should each one do?

Left reasons as follows: Top will either defect or remain silent. If he defects, then I do better by defecting: 9 years in jail instead of 10. If he remains silent, once again I do better by defecting: 0 years in jail instead of 1. Since in both circumstances I do better by defecting, that is my best strategy.

The game is symmetric in the two players, so Top's reasoning is exactly the same. So, evidently, each defects, and gets 9 years in jail. It's curious: if neither had defected, they would have only a year each. But the reasoning is impeccable.

We can represent the preferences by means of arrows

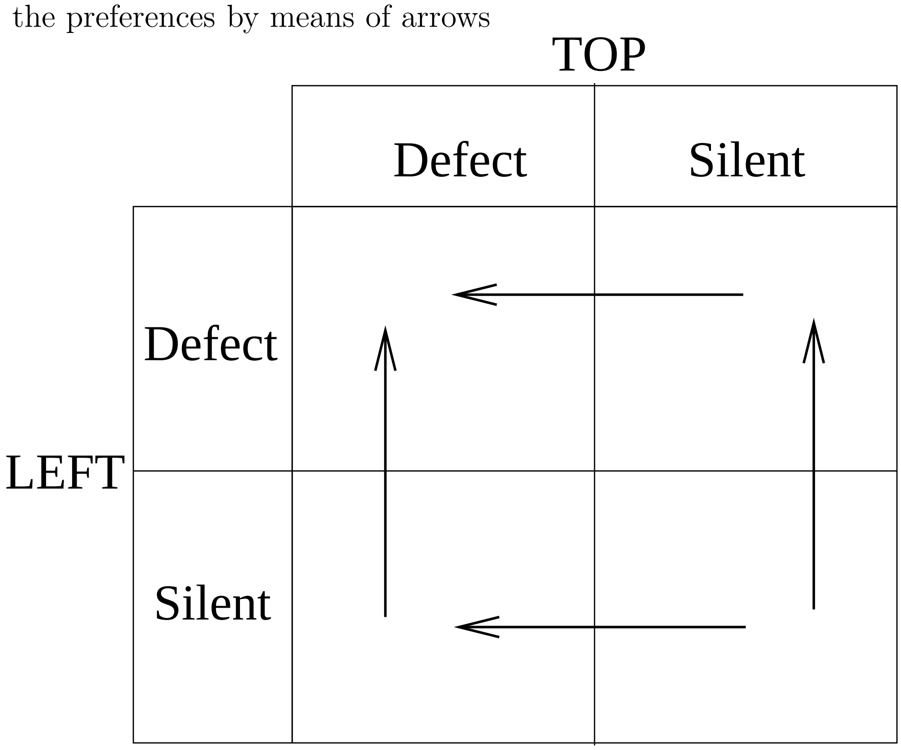

5cm

which may make it easier to follow the reasoning, and suggests the existence of a dynamic.

Here is a modified version of the game. It's called the *Diner's Dilemma*. Six people go out to dinner together, and agree to divide the bill equally. There are two choices, the expensive meal or the very expensive meal. What does each diner do? We assume they are all too polite to talk about money, so each reasons in private.

If I choose the expensive meal, I get all the benefit but only pay one sixth of the extra cost that my choice entails. Depending on how strongly I prefer the fancier meal, and how great is the difference in price, it may be that alone I would choose the cheaper meal, but in company I choose the more expensive. We can easily engineer this by assigning suitable values. Let

P be the pleasure of eating the very expensive meal

p be the pleasure of eating the less expensive meal

E be the cost of the very expensive meal

e be the cost of the less expensive meal.

Suppose, ordering them on a single scale, that

$$E \quad > \quad P \quad > \quad p \quad > \quad e.$$

My "payoff" is $P - E$ if I choose the very expensive meal, and $p - e$ if I choose the less expensive. As $P - E < 0$ and $p - e > 0$, it follows that if dining on my own, I would prefer the less expensive meal.

But dining in company, each diner's calculation is as follows: let R be the total cost of everyone else's meal. I have no control over that, of course. If I choose the less expensive meal, my benefit will be

$$p - \frac{R}{6} - \frac{e}{6}. \tag{0.1}$$

If I choose the more expensive meal, my benefit will be

$$P - \frac{R}{6} - \frac{E}{6}. \tag{0.2}$$

If

$$P - p > \frac{E}{6} - \frac{e}{6}$$

then (0.2) is greater than (0.1) and so acting on pure self-interest I choose the *more* expensive meal.

Again, the reasoning is the same for all the players, so all make the same decision, and *all end up eating, and paying for, the more expensive meal, against their previous wishes.*

Notice that the increase in eating pleasure is $P - p$; the increase in spending pain is $\frac{E-e}{6}$.

The reason that the accompanied diner's behaviour is different from that of the solitary diner is that the costs resulting from each decision are shared but the benefit resulting from each decision is for the individual alone. But curiously, instead of each diner benefitting from this increase in selfishness, all suffer.

Now suppose that the number of diners is not 6 but 6,706,993,152.

The key contribution to clarifying this state of affairs was made by the mathematician John Nash in the 1950's. It is the notion of the *Nash Equilibrium.* A choice of strategies (one for each player) is called a Nash equilibrium if for each player, his strategy brings him the best payoff *given the stategies of all of the others*. In other words, he cannot improve his payoff by a unilateral change of strategy. A game may have several Nash equilibria, or none. It is a famous theorem of Nash that in a purely competitive game, if we allow a "mixed strategy", in the form of a probability distribution, then there is at least one Nash equilibium. Put differently: if we imagine that the game is played a very large number of times, independently, and that the players strategy follows a probability distibution, i.e. do this $x\%$ of the time, do that $y\%$ of the time, etc, then there is a collection of probability distributions, one for each player,that together form a Nash equilibrium.

In any case, there are two main points I want to make about a Nash equilibrium. The first is that *this is what people tend to do*. They may not want to, but in the absence of cooperation, this is what they do do. Once the players in a purely competitive game have found a Nash equilibrium, either through a purely rational examination of the range of available stategies, or through an evolutionary process of trial and error or natural selection, then it can be rather hard to depart from it.

The second is that a Nash equilibrium may be extremely sub-optimal. In the two examples we've looked at, the players would get a much better outcome if they were able to discuss their options and agree to cooperate. Game theorists view this as changing the game itself. If my payoff in the new game is the sum of the payoffs of all the players in the old game, then my strategy will change to reflect that. This intellectual clarity has cost them something in reputation, though; their insistence that the prisoner's dilemma is not solved by cooperation (because given cooperation the game is no longer the prisoner's dilemma) has led many people to believe that game theory insists on a cynical and reductive view of human nature.

One of the ways in which the game may change is if it is played repeatedly, with players remembering the results of previous plays. One must be careful to distinguish between a finite iteration and an infinite iteration. In a finite iteration of the prisoner's dilemma then simple reasoning shows that nothing changes. At the last play, players have nothing to gain by good behaviour. Hence the optimal strategy is to defect. Since this is what will certainly happen in the last play, my behaviour in the penultimate play will have no effect on your behaviour in the last, so once again my best strategy is to defect. And so on back to the start, by induction.

On the other hand, if we allow an infinite iteration (or, more realistically, an iteration in which we do not know which is the last), then behaviour can change dramatically. If my behaviour this play *can* influence your behaviour in the next, then maybe I should change my ways. This was investigated in a striking way by Robert Axelrod in a book called *The Evolution of Cooperation*, published in 1984 ([1]). His introduction opens:

> When should a person cooperate, and when should a person be selfish, in an ongoing interaction with another person? Should a friend keep providing favours to another friend who never reciprocates? Should a business provide promopt service to another business that is about to be bankrupt? How intensely should the United States try to punish the Soviet Union for a particular hostile act, and what pattern of behaviour can the United States use to best elicit cooperative

behaviour from the Soviet Union?

Axelrod invited experts in game theory to submit computer programmes containing a strategy for an indefinitely iterated prisoners' dilemma (that is, which took no account of the fact that only finitely many iterations would actually be played). Programmes played against one another in single combat over 200 iterations, and Axelrod kept a cumulative total of the payoffs and analysed the results in his book. Each programme also played its own twin, and against a programme which selected its plays randomly. There were 14 entries, from game theorists in economics, psychology, sociology, political science and mathematics. Some of the programmes were extremely ingenious and complex. Some were "nice", some were "nasty". The winner was the simplest, TIT FOR TAT, submitted by Anatol Rapaport. In TFT, behaviour is governed by a simple maxim: if you defect at iteration k then I defect at $k+1$. If you cooperate at k, I cooperate at $k+1$. I always open by cooperating.

Axelrod defined a nice programme as one which is never the first to defect. He found, rather encouragingly, that the top scoring eight programmes were all nice.

I will not go into more detail about his experimental results. The book is easy reading for a mathematician, and leaves you feeling hopeful: cooperation pays, and can be promoted by simple legislation, even among computer programmes. It can even thrive, and produce good results, in circumstances one would imagine to be extremely unfavourable. The introduction, by Richard Dawkins, concludes

> The world's leaders shold all be locked up with this book and not released until they have read it. This would be a pleasure to them, and might save the rest of us. *The Evolution of Cooperation* desreves to replace the Gideon Bible.

In fact we learn from his introduciton that Dawkins, notorious for his notion of the Selfish Gene, was one of those initially invited to play Axelrod's game, but didn't get round to submitting a programme.

Back to climate change. We saw while discussing the diner's dilemma that the sub-optimal Nash equilibrium becomes more pressing, more attractive, more seductive, the larger is the number of diners. This is our situation today, facing the exhaustion of a number of the world's different resources. The term "The Tragedy of the Commons" was coined by the ecologist Garrett Hardin in his essay of that name ([5]) in 1968 (still available on the internet). Much of rural England was common land until the eighteenth century, and even on into the nineteenth. On this common land, people could graze their livestock and harvest firewood as they wished. But beginning in the sixteenth century, Parliament passed laws (Inclsures Acts) allowing landowners to fence off and take possesssion of tracts of common land. The injustice of this land grab inspired the rhyme

> We hang the man and flog the woman
> Who steals a goose from off the common
> But let the greater villain loose
> Who steals the common from the goose.

But though this pulls at our heart strings, the arguments in favour of the Inclosures Acts were not without merit. The arithmetic of profit and loss from the exploitation of common land is exactly the same as in the diner's dilemma: the cost of my sheep grazing the common land is shared among all the users, but I alone profit from the wool and mutton. The result is over-grazing and low productivity. Marxist historians point to the need to displace peasants from the countryside in order to form an urban proletariat to work in the factories, and I am not enough of a historian to know how much relative weight to give to the two explanations. But I would say that it is a traditional human mistake to attribute suffering and disasters to the wickedness of knowing agents rather than to the unwitting stimuli of Nash equilibria.

Hardin's interest was not historical though, but ecological and political, and his title is now used in game theory and economics to refer to the tendency to destroy shared resources through overexploitation for private gain.

Let me even the score of my attacks to left and right by reading what he said, in 1968, about classical free-market economics, at that point set for a a great come-back though the work of Milton Friedman.

> We can make little progress in working toward optimum population size until we explicitly exorcise the spirit of Adam Smith in the field of practical demography. In economic affairs, The Wealth of Nations (1776) popularized the "invisible hand," the idea that an individual who "intends only his own gain," is, as it were, "led by an invisible hand to promotethe public interest." Adam Smith did not assert that this was invariably true, and perhaps neither did any of his followers. But he contributed to a dominant tendency of thought that has ever since interfered with positive action based on rational analysis, namely, the tendency to assume that decisions reached individually will, in fact, be the best decisions for an entire society. If this assumption is correct it justifies the continuance of our present policy of laissez faire in reproduction. If it is correct we can assume that men will control their individual fecundity so as to produce the optimum population. If the assumption is not correct, we need to reexamine our individual freedoms to see which ones are defensible.

What should mathematicians do about any of this? Well, first and foremost, we are not mathematicians. We are citizens, of the global city. We all have an interest in preserving the global resources which make our lives possible. From this I deduce that it is, for example, legitimate for me, who is not an expert on game theory, climate modelling or economics, to try to promote activity on these fronts, if it seems to me that not enough is being done. And I would say that this is pretty clear. Not enough is being done. If anything, the credit crunch should bring home to us that our elders and betters don't really have any idea of what they are doing, and may well be in the grip of Nash equilibria more ferocious and just as sub-optimal as the ones I described above in the toy games.

Our politicians, subject as they are to periodic electoral challenges, have a strong electoral disincentive against undertaking policies which seriously address the problems of climate change. The reason is the same as always: the cost of alleviation is borne by the alleviator, while the benefit is distributed among everyone in the world and their descendents. Electorates tend to have little empathy with everyone in the world and their descendents, and demand rising living standards here now. The problem is epistemological.

Our first job: understand what is going on. We are well qualified. In almost all disciplines bearing on the issues I'm discussing today, the hard part is the mathematics. This we can handle.

At one extreme of the spectrum of activity open to us as mathematicians is climate science itself, in which there are still enormous mathematical challenges, of a very modern kind. Even though the fact of climate change and global warming is no longer in dispute, its precise degree, and its effects on rainfall, agriculture and vegetation are very uncertain. The climate models with which we hope to answer these questions depend on very large scale computation. In particular, on the skilful assimilation of the results of models into the collection and interpretation of data. This brings together a number of areas of applied mathematics, in which mathematicians in this department are working. Warwick has a nationally leading Centre for Scientific Computing which is in fact still somewhat underused.

At the other extreme is political economy: the design and evaluation of economic and legislative mechanisms to reduce greenhouse gas emissions and promote clean technology. Last week the Mathematical Sciences Research Institute in Berkeley California ran a meeting on Game Theory and Climate Change Treaties. The principal organiser was Chris Jones, who has a joint appoinment here in Warwick and in Chapel Hill North Carolina. It was attended by a diverse audience of mathematicians, economists and in-betweens.

This brings me to my own project, which I will describe briefly. Economists have estimated the cost of climate change abatement as rather small. The Stern review [10], for example, which is one of the most authoritative, reached the following conclusions:

> Using the results frm formal economic models, the review estimates that if we don't act, the overall costs and risks of climate change will be equivalent to losing at least 5% of global GDP each year, now and forever. If a wider range of risks and impacts is taken into account, the estimates of damage could rise to 20% of GDP or more.
>
> In contrast, the costs of action – reducing greenhouse gas emissions to avoid the worst inmpacts of climate change – can be limited to around 1% of global GDP each year.

One of the speakers at the Berkely meeting was Robert Pindyck, Bank of Tokyo-Mitsubishi Professor of Economics and Finance at the Massachussetts Institute of Technology, MIT. His contribution was an analysis of the willingness of society to pay for climate change abatement.

From the introduction to his paper ([9]):

> I estimate the fraction of consumption $w(\tau)$ that society would be willing to sacrifice to ensure that any increase in temperature at a future point is limited to τ. Using information on the distributions for temperature change and economic impact from studies assembled by the IPCC and from integrated assessment

> models (IAMs), I fit displaced gamma distributions for these variables. Unlike existing IAMs, I model economic impact as a relationship between temperature change and the growth rate of GDP as opposed to the level of GDP. This allows warming to have a permanent impact on future GDP. I find that the fitted distributions for temperature change and economic impact yield values of $w(\tau)$ above 2% or 3% for small values of τ only for extreme parameter values and/or substantial shifts in the temperature distribution – which does not support the immediate adoption of a stringent abatement policy.

In response to a question about the level of defense spending in the US during WWII (the answer was roughly 20% of GDP) he made clear his own view that the threat of, for example, nuclear terrorism is more serious than the threat of climate change.

I find this response, and the conclusions of his and Nicholas Stern's reviews, and others along similar lines, very hard to square with what I as a layman piece together from a lot of rather diverse sources about the level of threat from climate change. James Lovelock predicts human survival only in small regions around the north and outh poles, with the intervening regions converted to desert.

So what is going on with these economists? Alternatively, what is going on with us? We seem to be two groups of reasonably well informed and generously motivated individuals with radically differing estimates of how much needs to be done. What gives rise to these differences? Is our ignorance of economics a convincing reason to leave the job to the economists? I do not think so. I am encouraged in my thinking not just by a desire for self preservation, which can justify practically any act of intellectual arrogance or madness, but also by the belief that economists are excessively dominated by numbers. For us mathematicians, on the other hand, numbers are only a small part of our repertoire. We know how to think about knots, smoke rings, epidemics, evolutionary dynamics, simplicity and complexity in ways that steadily expand outwards from an initial core of numbers.

If it is us who are right, but the economists who are listened to by governments, then we are in trouble. If the economists are right, and our anxious estimates are way too high, then the worst that will happen is that we will have wasted a lot of time. But that's another thing mathematicians know how to put up with.

In his book *Blink* ([4]) Malcolm Gladwell shows how snap judgements under some circumstances give a better answer than the work of learned professors. In most of his examples, it is the snap judegements of experts that is superior to the considered and mediated judgements of other experts, but nevertheless his book encourages one to take one's intuitions seriously.

I end with a description of another toy game that I learned of from a Portuguese mathematician called Jorge Pacheco, who gave a talk here in Warwick two weeks ago (see [8]). It's called Save the Planet, and is played as follows:

There are 6 players. The game is played over 10 rounds. Each player is given 40 euros to start with. In each round he makes a contribution from his 40 euros to a fund to save the planet, of

0 Euros (selfish)   2 euros (fair)   or   4 euros (altruistic).

To save the planet, the fund must reach 120 euros after the 10 rounds. This means an average of 2 euros per player per round - the fair contribution. If this total is achieved, the players get to keep the money they haven't spent. If the total is not achieved, then with 90% probability the planet is destroyed – the players have to hand back all the money and go home empty handed. There is a 10% chance that the climate predictions are wrong and the players can keep the money they've retained, even if the 120 euro total has not been reached.

So what happens? Pacheco had some experimental data, which you can find on the slides of his talk, available from our departmental website. The idea of the game originates in an article by four members of the Max Planck Institutes of Limnology and Meteorology in Germany. They studied several variants in which the size of contributions was made public, or not, and players got to know one another through other games. It seems they also told the undergraduates playing the game that the fund would contribute to the cost of a newspaper advertisement publicising the dangers of climate change. The four authors claimed to be encouraged by the results, published in the Proceedings of the (US) National Academy of Sciences in 2006 ([6]).

But I think it's more interesting for you to play games than for me to tell you about them. Of course games are not mathematics, but understanding the behaviour of their participants is a combination of psychology and mathematics that may yet have a significant role to play in averting the worst effects of climate change.

References

[1] Robert Axelrod, *The Evolution of Cooperation*, revised edition, Basic Books, 2006

[2] Ken Binmore, *Game Theory, A very short introduction*, Oxford University Press, 2007

[3] Freeman Dyson, The Question of Global Warming, (review of [7]), New York Review of Books, June 2008, available online.

[4] Malcolm Gladwell, *Blink*, this edition Back Bay Books, 2007

[5] Garrett Hardin, The Tragedy of the Commons, Science, 162(1968):1243-1248.

[6] Manfred Milinski, Dirk Semmann, Hans-Jurgen Krambeck, Jochem Marotzke, Stabilizing the Earths climate is not a losing game: Supporting evidence from public goods experiments, Proceedings Nat. Acad Sci. USA, vol 103 no. 11, March 2006, 3994-3998

[7] William Nordhaus, A Question of Balance: Weighing the Options on Global Warming Policies, Yale University Press, 2008

[8] J. Pacheco, talk at Warwick Mathematics Institute MIR@W Day, April 27, 2009, slides available at http://www2.warwick.ac.uk/fac/sci/maths/research/miraw/days/

[9] Robert Pindyck, Uncertainty, Extreme Outcomes, and Climate Change Policy, Working paper, Spring 2009, available at http://web.mit.edu/rpindyck/www/papers.htm

[10] N.Stern, *The Economics of Climate Change, the Stern Review*, Cambridge University Press, 2006