Chapter 3: Question and Answer

MEMBER OF THE AUDIENCE:

This is a question of clarification for Chris Hope about the treatment of uncertainty. As I understand it, the probability distributions that you are drawing for outcomes start from a range of values for each of the parameters that go into the model. So you have a joint distribution of values for the parameters, and depending upon how the joint distributions interact in the model, you generate the distribution of outcomes. So you are not working with uncertainty per se, with the distribution of outcomes in a probabilistic sense, but rather with the distribution of simulations, if I understand you correctly.

It's true that you don't have enough information to put together a joint distribution, but why not start with some prior distribution and, as the information comes in, update the prior distribution and see what sort of outcome probability a distribution of outcomes generates?

CHRIS HOPE:

For each of the different parameters we try to define uncertainty ranges that reflect those in the literature or in research. That's why we have a range of one-and-a-half to five degrees for climate sensitivity, for example. Then we run these ranges as joint probability distributions all the way through the model. From 10,000 runs of the model, we pick from each of these joint probability distributions for the inputs and build up a probability distribution for the output.

Most of the time you want to design your models so that the inputs are independent of one another. But sometimes you have to deal with input correlations. You can the PAGE model with correlated input distributions and work out what the output distribution is.

And, as my last slides showed, you can then talk about what happens if you get better information at some point in the future and you do some hedging strategy up to there, and then when we get better information, we split the strategies up according to that information. That can tell you something about how much that information is worth to you. Those slides only showed the value of information for a scientific parameter—climate sensitivity—but obviously, it would be nice to also model the value of extra information for things like the impact function exponent.

MEMBER OF THE AUDIENCE:

How useful, and significant in your analysis, are policies that accelerate the demographic transition?

NICHOLAS STERN:

Policies that accelerate the demographic transition are valuable in their own right. I think we understand that what helps drive the transition is policies that make good sense to follow anyway: promotion of human rights, promotion of economic development, job opportunities for women, education for girls and women, access to reproductive health services. Those are the kinds of things we know from looking at the way in which fertility rates have dropped, say, in Bangladesh from six to three over the last 25 or 30 years.

Such policies would help, of course, to mitigate global warming, but I wouldn't particularly want to go that route in arguing for them: I'd want to argue for them in their own right.

MEMBER OF THE AUDIENCE:

This is another clarifying question. How reduced-form is your model? Do you actually try to trace out the future of energy consumption and pinpoint when the transition from oil to coal is taking place, and so on? Do you estimate how much the temperature is going to rise on the basis of exactly which technology is being used at which time? Or do you just make an assumption about how fast the temperature is going to rise and an assumption about how costly it is if the temperature rises?

CHRIS HOPE:

It's a very reduced form model. We do not explicitly model the energy sectors, the use of the different fuels, and so on; we start at the point of saying what the emissions of the different greenhouse gases are. So, we need to feed into the model what the emissions are predicted to be of CO_2, methane, sulfurhexafluoride, and the other gases that are important in causing climate change. From there, we can work out what's going to happen to the concentrations, the forcing, the temperature, all the while bringing in the chain of uncertainties that Sir Nicholas has been talking about, and work through to the impacts at the end and the uncertainties there.

So, PAGE doesn't have an explicit energy model on the front of it. It would be nice to have one. But since the Inter-Governmental Panel on Climate Change (IPCC) has already done a lot of work to build up the emission pathway, and provided this readymade for us to use, I decided that making our own explicit energy model wasn't the best way of using our effort.

NICHOLAS STERN:

On the cost side, in thinking about how to start to control the emissions and get them to peak in 20 years time and gradually come down, we did look at the implications of different kinds of technologies and when they might come into use. In that kind of analysis you do get the different balance of energy use from different technologies, but it's not integrated into the full model.

My own view on these models is that at some point you've got to try to stop loading lots of things in. On the whole we looked at the cost side separately from the overall modeling, to allow us to get involved in much more detail than could be handled in the full model.

MEMBER OF THE AUDIENCE:

Regarding the PAGE model, I'm curious about the choice of the A2 scenario as used in the Stern Review. I know that a number of the assumptions used in this scenario—particularly about population but also about economic growth rates—may become somewhat problematic, especially if you extend the model out to the year 2200. I'm curious if you've tried running the PAGE model with, say, the A1F or A1B scenarios and if so, how that changes the outcomes.

CHRIS HOPE:

One reason why we used the A2 scenario was that this is one of only two scenarios in the third IPCC report that has some probabilistic information on it. There were several runs in the report which you could use to understand the range of answers that might come out if you use that scenario.. The other scenario that the third IPCC report described in that kind of detail was scenario B2, and I've done quite a lot of runs with that as well. A third scenario, elaborated since the third IPCC report, is the common POLES-IMAGE scenario: business as usual.

You can run whatever you like as your business-as-usual scenario, add some assumptions about uncertainty, and then see what impact that has on the answer. One of our interesting results is that the social cost of carbon doesn't vary much depending on which scenario you impose it on.

NICHOLAS STERN:

The assumption made about population does matter for the balanced growth equivalent. There is quite a heavy population load in the A2-1 scenario. So we did some back-of-envelope sensitivity analysis to see what difference working with smaller population numbers might make, given the same kinds of emissions. This showed that with smaller population, the environmental damage is somewhat less.

MEMBER OF THE AUDIENCE:

An issue for policymakers is whether the appropriate social cost of carbon to use in policymaking is the business-as-usual social cost of carbon or the significantly lower social cost of carbon that is associated with a stabilization trajectory to 550 parts per million (ppm) of CO_2 equivalent? What are your thoughts on that?

NICHOLAS STERN:

Any shout-out price, which is a marginal concept of course, is problematic in a modeling structure where you're dealing with (in the model itself) very big non-marginal effects and if further you're dealing with policies which themselves could make very big differences. This means that you have to be very careful to attach to any estimate of the social cost of carbon a statement about the kind of path to which it corresponds.

Here Chris Hope was very explicit. His social cost of carbon attaches to the business-as-usual path. So, suppose you take the business-as-usual path with these rising emissions and you perturb it a little bit at a given point in time. What happens to the subsequent path, and then the difference between the two, is the damage caused by the extra carbon.

With a path that is more sensible than business as usual, and entails controlling emissions into the future to stabilize at 550 ppm, the social cost of carbon will be lower because the stocks of emissions over time will be lower.

This is a classic example from applied cost/benefit analysis, of needing to be explicit about the relationship between the shout-out price and the overall path that is being followed. It does, of course, create a difficulty for the policymaker: suppose for example that you're assessing the social cost of carbon in the context of a road-building project that will save some travel time and some fuel. What price do you use for the carbon associated with the fuel that is saved? If you think the world is going to be sensible about future emissions, you'll choose a lower price. If you think the world is not going to be so sensible, you'll choose a higher price. It seems to me that if the carbon price really matters to the investment decision that's being taken, you have to look at it very hard and see what kind of probabilities attach to the adoption of different kinds of policies.

I wouldn't place huge emphasis on specific estimates of the social cost of carbon, given that these vary widely. The concept is a valuable one in analysis but much less so as a guide to policy. That's why in my presentation I didn't lay a lot of emphasis on the social cost of carbon as a policy tool, and why we argued that, from the point of view of the economics of risk of the Marty Weitzman kind, one should set stabilization goals and find the path associated with them.

Any given stabilization goal will have a corridor of paths associated with it, and you use your price mechanisms to decarbonize within that corridor of paths, but revise your decisions from time to time as better information comes in. And in that revision, a field for the social cost of carbon would be helpful, but we didn't set that up as the central driving force of the decarbonization policy story as we described it.

CHRIS HOPE:

Nick is very well aware of the political process. I tend to think of the social cost of carbon as what you should use to set any carbon tax that you might have, and I would agree with Nick that you should probably look at what the carbon tax should be under an optimal, rather than under a business-as-usual, path of emissions.

I actually think that you should probably set your carbon tax a little bit higher than the mean social cost of carbon—for various reasons to do with the things that have been missed out of the calculations, but also because the carbon tax is going to be replacing other taxes that distort decision making in the economy.

Even if you accept this point of view, it's still an open research question as to how much the social cost of carbon will change if we move from something like a business-as-usual scenario to something like an optimal path of emissions over time. Maybe the cost will drop by a factor of two or three, or maybe it will not drop much at all. The answer seems to depend quite a lot on your assumptions about what's going to happen to other gases, like sulfates and other background gases in the atmosphere, as you move to your optimal path of emissions.

MEMBER OF THE AUDIENCE:

Was acidification of the oceans included in your calculations of social cost?

CHRIS HOPE:

The PAGE model, as run for this kind of study, doesn't take into account the other impacts that there might be of fossil fuel use. So neither does it explicitly take into account the co-benefits from reducing fossil fuel use, which include the reduction of acid rain.

It is possible to run the model along with other models that track those things explicitly, and find out the total benefit of cutting back, let's say, emissions from fossil fuel. But you have to be quite careful how you do that: if, for example, some of your policy action is to reduce the emissions of CO_2 by reducing deforestation, this won't necessarily yield the same kinds of co-benefits as from reducing the burning of fossil fuels.

NICHOLAS STERN:

I think this is a good example of why you'd want to go for disaggregated descriptions. In our report we did look across a whole range of these kinds of problems, and the decision challenge, then, is: Would you pay one %of GDP to drastically reduce this whole collection of types of risk?

MEMBER OF THE AUDIENCE:

On climate sensitivity, the IPCC process takes models from all the countries in the world that are able to submit data. There is really no quality control—there is no way to check. But if you take the view that those models that simulate current climates well should be considered the more reliable ones for forecasting the future climate, you find out that the climate sensitivity is at the highest range of the IPCC estimates. This is something that the IPCC has decided to investigate further, starting from its next assessment.

NICHOLAS STERN:

I think that's an example of the ways in which we were cautious. We didn't know how to incorporate those kinds of possibilities, though there are some strong signs in the IPCC report. I do think there are lots of uncertainties out there and risks of a major kind that are simply not in the story as we told it.

CHRIS HOPE:

I have done some runs with the higher numbers that seem to be coming through from Murphy and Stainforth and so on. These numbers seem to increase the mean values for the social cost of carbon by at least 50 percent, and possibly a bit more. This is exactly the benefit of this kind of model: when you get better information, you can find out what effect it has on the policy-relevant numbers.