Chapter 18

Controlling climate change economically

James Mirrlees

James Mirrlees was born in Minnigaff, Scotland, in 1936. He studied mathematics at Edinburgh and Cambridge Universities, and received a PhD in Economics at Cambridge University in 1963. He has been Professor in Oxford and Cambridge, where he is Emeritus Professor of Political Economy, and is now also Distinguished Professor-at-large in the Chinese University of Hong Kong. In 1996 he was awarded the Sveriges Riksbank Nobel Memorial Prize in Economic Sciences 'for his contributions to the economic theory of incentives under asymmetric information'. Mirrlees had focused on the design of income taxation but soon realized that his method could be applied to many other similar problems where decision-makers are faced with asymmetrical or incomplete economic information. Up to the present day, optimal public economic policy remains his main research interest, as well as development economics.

Note: This chapter is a commentary on chapter 17.

When people do things that emit greenhouse gases into the atmosphere, they cause damage – damage to other people; in the present, in the future, and all over the world. To an economist, it is clear that emitters should pay a price equal to the value of the damage caused. That is what we call the carbon price (though non-carbon greenhouse gas emitters need to be charged too). With carbon prices in place, now and in the future, people will burn fossil fuels only if the value of doing so is greater than the cost of the fuel and the combustion device, and the damage caused by the increases in temperature it brings about. Supplying energy by less damaging processes will pay off. The invention and development of new energy technologies will be more profitable. However, leaving aside general principles it is not easy to estimate the damage caused by global warming, and there is considerable disagreement about what should count as damage, and about the economic cost of that damage.

As many non-economists have recognized, our relatively carbon-free atmosphere is an exhaustible resource. This is the approach taken by Edenhofer and his co-authors who, instead of computing the social cost of carbon, consider the simplified problem of optimally allocating a maximum admissible quantity of emissions – a carbon budget. According to Hotelling's principle, the price of an exhaustible resource should, uncertainties apart, rise at a rate equal to the rate of interest. Consequently, Edenhofer and his co-authors report their finding that the carbon price should rise over time, at least for an initial period until backstop technologies become more competitive. But if global warming is already causing damage, and going to cause more in the immediate future, Hotelling's principle has to be modified. I am sure that is the case. It is quite possible that the carbon price should not be rising; it should perhaps already be as high as it is ever going to be. In particular, rough calculations suggest that it should already be much higher than the prices that have so far emerged in cap-and-trade markets.

Carbon prices can be determined in markets, in order to equate the emissions people wish to make to the quantity of emissions that is compatible with mitigation goals. Many have argued that it is easier to estimate the desired quantity of carbon emissions than to estimate the value of the damage caused by emissions; others that it is easier to estimate the carbon price. Neither position is tenable. If we seek an optimal solution, prices and quantities have to be estimated together, as I shall explain. It is possible to develop a plan for present and future emission levels that would have a good chance of keeping damage from warming within tolerable bounds. Such a plan can be resolved into target quantities for individual countries or industries, as in the case of the Kyoto Protocol. While proceeding in this way has the advantage that it is readily understood and relatively easy to discuss, it is, however, unlikely to gain the universal coverage necessary to achieve the desired effect.

Even taking the future course of emissions as a given, it is a daunting task to estimate the damage. Not only do we need to estimate the extent to which a unit of

emissions will reduce national incomes (the total of the individual incomes that we actually care about), we need to do that for long periods of time, far into the future. We also need to assess forms of damage that are not easily quantifiable in monetary terms, particularly loss of life and environmental destruction. Certainly, future economic damage should be discounted. Equating the discount rate to the average rate of return to capital investment may be helpful in estimating it; but that rate of return cannot be estimated independently of considerations of future climate change damages. In any case, the discount rate should not be applied to the value of life and environment lost. That matters, because it means that the equivalent value of such destruction is important now, even if it happens a century hence.

There is also a strong case for increasing damage estimates to allow for uncertainty about the future size of national incomes, which will partially offset the discount rate. Anticipated lower national incomes in the future would require us to adjust the discount rate downwards. That is another reason why damage must be estimated a long way ahead. I agree with Nicholas Stern and Martin Weitzman that the effective discount rate for standard economic damage should be low, perhaps two or three percent, though I think we all have different reasons.

The level of damage at any future time depends on how much global warming actually takes place. Our estimate of the carbon price therefore also depends on the extent of global warming, which will be influenced by the controls and carbon markets and taxes that are put in place. Certainly very different levels of damage are possible, up to widespread loss of life. The question is whether marginal (i.e., additional) damage increases with the level of damage.

It might be thought that, following the initial decades in which global warming begins to have a significant impact (i.e., the present day), the marginal damage from further unit increases in emissions will not vary much with the level of future damage. This view is supported by estimates quoted in the Stern Report, which imply that the marginal damage will not increase from 2 °C to 4 °C of warming. These figures may well reflect the simple ways of estimating future damage that have so far been employed in attempts to quantify the effects of warming. However, when one considers urban reconstruction, population displacement, and likely death tolls from high and rising water levels, all costs and damages that should not be discounted, it is hard not to believe that marginal damage will increase with the level of future carbon concentration, and will be very high if concentrations reach a level where the temperature increase is 4 °C or more.

If it is true that marginal damage from global warming will increase with the extent of warming, serious consequences follow for countries that wish to effectively address the climate change problem. If it becomes clear that many countries are not going to radically reduce their emissions, the marginal social cost of warming will be all the greater. It might then be the duty of compliant countries to reduce

their emissions to zero as quickly as possible. The value for them of a global deal is great, as are the advantages of subsidising carbon savings in other countries.

The claim that marginal damage increases with carbon concentration provides an answer to the following question: how can we be sure that a carbon price calculated by estimating future marginal damage will be high enough to assure effective mitigation that keeps the carbon concentration from pushing temperatures above a tolerable level, say 3 °C? The price estimate based on future concentrations above that level would surely be more than high enough to prevent it. Effective mitigation measures that draw down carbon concentrations would surely be undertaken since the cost of radical cuts in greenhouse gas emissions is comparatively low (it is reliably estimated to be only a modest fraction of national incomes, less than 5%). This argument also shows that we cannot estimate the appropriate carbon price simply by discounting future marginal damage, since that depends on what future carbon prices will be. We need a full dynamic model, to be solved simultaneously for optimal emissions and optimal carbon price. A cruder version of this model is to estimate the desirable carbon price on the basis of realistic forecasts of temperatures and water levels, given policy commitments as shown by actual carbon prices, nation by nation, and adjust it year by year. I expect that the desirable carbon price would fall over time, as countries improved their mitigation policies.

Is it easier, then, to estimate the quantity of emissions that we should have? It is easier to propose a particular quantity plan than a price-and-tax plan, and that is what the global-warming community has done, and what most governments have accepted, in a rather quixotic way. Laying out a carefully estimated scenario is indeed a great achievement, and makes some kind of international agreement possible. But it is something very different from the calculation of an optimal plan, which (some) economists might prefer. It leaves open the question of how much pollution each nation or person may cause. Now that is certainly hard to estimate, in the sense of giving an optimal prescription to each nation. The nearer decisions come down to the individual level, the more we must move to price rather than quantity, because different individuals should have different 'carbon footprints' – they have different tastes, live in places that create different needs, and have different possibilities for emission reductions.

The way that carbon markets work now creates many anomalies. The main one is that the level of carbon prices seems much too low, when compared with the sorts of figures that are estimated on the economic basis I outlined. The carbon price needs to be estimated on the basis of future damage, as a check on the level and effectiveness of the quantity constraints that carbon markets are supposed to embody. One reason for the low price is the provision for buying emission permits from emission-reduction projects in other countries. Of course that has the beneficial effect of reducing net emissions in other countries. But it is offset by increasing net

emissions above what was supposed to be the agreed level for the country, or group of countries, operating the carbon market. There is no net advantage for global warming from that. This flaw in the current implementation of carbon markets could have been remedied if the quantity of permits made available each year had been reduced by the amount expected to be purchased from emission-reducers. As things are, these external purchases brought the carbon price down, reducing the price to users of the kinds of consumption and investment that have a relatively large impact on global warming, and reducing the incentive to introduce and develop green technologies.

I am not arguing that there should be no subsidy for emission-reducing projects, such as forests. On the contrary, it is most desirable to encourage projects that absorb greenhouse gases. We should talk of a carbon price, not a carbon tax, because it should operate as a subsidy to activities that absorb greenhouse gases from the atmosphere, as well as a tax on the emission of greenhouse gases. The low carbon prices in the cap-and-trade markets are evidence that, for the time being, expansion of forests (to take a major example) and some other flora is the most efficient way of reducing global warming. Since the carbon price ought to be, based on any estimate, considerably higher than the current price, there should be a very rapid expansion of carbon-absorbing flora. But, for obvious reasons, it cannot last for long, because there is only limited space before the value of producing food makes further expansion undesirable.

The more general problem with carbon markets is inclusiveness. Not everything can be covered by a carbon market, though it seems governments are very far from requiring permits wherever they should. Farmers should have to purchase annual permits for their farm animals, just as they should be able to sell annual permits on the basis of their woodland. It is a complicated issue, however: the quantity of permits should be related to the type of animal, its age, and other characteristics. This will no doubt seem quite impossible in the European Union. Can we envisage it in India and Africa? Yet, if not, how are we going to get emissions down to 80% of 1990 levels by 2050, with cuts increasing in intervening years; or even the somewhat more modest ambition of the G8 countries? In many cases, it is simply easier to apply the carbon price as a tax than to require purchase of permits. The tax rate could be based automatically on the price in the carbon market. One argument for setting quantities of carbon permits (within each country) and allowing them to be traded in a market is that allowable emissions are then produced by those for whom it is most valuable. If the market does not cover the full range of emission activities, this will not happen to the extent it should.