12 Reflections for public policy: a drawing together and drawing apart. Comments on proceedings

Geoffrey Brennan
Australian National University

12.1 The great divide

My task this afternoon is to provide a kind of grand synthesis of all that has gone on in the course of this event. It is, I'm afraid, a task that is beyond me.

Indeed, instead of drawing the ends together, I am actually going to try to draw them apart. That is my ambition because — although everything on the table since Warwick McKibbin began to talk last night has been connected in one way or another with the 'environment' — I think that that commonality of subject-matter has served to mask what I regard as a more basic, and quite crucial *distinction (*one that I regard as ultimately economic in character) between the intranational and the international aspects of environmental policy. Simply put, there is I believe a categorical distinction between the 'global warming/climate change/greenhouse gas' issue, and everything else on today's agenda.

So, on the global side of my conceptual divide, I place the McKibbin and Stavins papers, with perhaps Suzi Kerr having a little each way. All of the other papers fall on the national and sub-national side of the divide. (As I shall try to argue, I consider that balance entirely proper.) So, problems of water extraction in the Murray–Darling system (the Buchan and Craik–Cleaver papers) or the NAP and NHT initiatives (Pannell) or the fine-tuning of political institutions and geographical domain in sub-national environments (David Brunckhorst's paper) are all examples. I place Henry Ergas' paper in this group because it is concerned primarily with the purely domestic issue of how best to use the revenue from carbon-reduction schemes.

Equally, the Libecap and Freebairn papers (though at a more general level) are directed at the national policy design problem — dealing with the comparative evaluation of alternative instruments for externality problems *within* a given political boundary. Although both papers have an 'environmental' gloss, the basic issues are ones that arise in a wide variety of policy areas: the exercise is largely one of taking a general analytic framework from public economics associated with Samuelson and from micro-economics associated with Coase (and perhaps political science associated with Elinor Ostrom) and applying that framework in the environmental setting.

Of course, the two dimensions of environmental policy — the national and the global — are related, in a variety of ways and at a variety of levels. And I shall later want to say a little about a number of those connections. But at this point, I want to underline the logic of the global/national distinction — partly because I am deeply uneasy about much of the discussion of carbon emissions policy. I am uneasy because I am not quite sure whether people who ought to know better are putting their blind eye to the telescope as a kind of Nelsonian heroic gesture or a rhetorical manoeuvre; or whether they don't really understand the magnitude of the challenge, or worse, the real nature of the problem.

Let me put the point a slightly different way. I confess that I have never been much impressed with that economist joke that has as its punch line: 'Assume a can-opener!'[1] It has never seemed to me to be either particularly funny or particularly apt. Until recently. Because these days I hear a lot of talk in exalted policy circles about the details of Australia's projected policy to combat CO_2 emissions and stop global warming — and I can't help wondering where the can-opener has come from!

So, back to basics. There is a long tradition of thought, beginning perhaps with Thomas Hobbes, running through David Hume and Adam Smith and familiar to economists from Samuelson's canonical public goods papers, that sees the primary rationale for government as lying in what we would identify as the n-person version of Tucker's prisoner's dilemma. The common claim is that we need institutions that possess coercive power (the power specifically to tax and/or regulate) because decentralised decision making cannot produce public goods in anything like optimal quantities. As Bob Stavins rightly remarked, we cannot rely on voluntary behaviour to solve the policy problems that our emissions create. We just wouldn't be here talking about global warming or climate change issues if that were so.

[1] I assume that every economist knows that joke. Those who do not can find it on Google under 'economist jokes'. Some sources attribute it to Paul Samuelson.

This is something that economists *know*. We know it because we understand the structure of individual interaction and we have good evidence for thinking that people are less than totally altruistic.

Global carbon emissions reductions are, as I understand the science, a classic case of a global public good. That is, the greenhouse gases that are released from any source are pretty rapidly diffused into the atmosphere — if Australia's carbon emissions affected only the Australian climate, then things would be different. But the effect on Australia's atmosphere of any reduction in Australian carbon emissions is just to reduce carbon by the proportion represented by Australia's share of world carbon emissions. Australia's actions on CO_2 emissions therefore have negligible effect on Australia's climate; Australian policy to reduce carbon emissions will only have a proportionate effect on Australia's climate if Australia's policy action *causes* all other countries to impose the same reduction in CO_2 emissions — and it is the essence of independent action that Australia's action does *not* cause identical policy in all other countries!

Recall that in an independent adjustment equilibrium in the provision of a public good, it must be the case for every individual contributor that the marginal value of a dollar's contribution to public goods provision is a dollar. (At optimality, each, on average, has a marginal evaluation of the public good equal to 1/n dollars; so it won't be rational for any individual to contribute at that optimal level unless she is part of an n-person compact in which each contribution is matched by the contributions of all others.) If one or other of the individuals in question is highly altruistic, then that individual may contribute a large amount to the public good. But an increase in that individual's contribution is, under a variety of assumptions, likely to cause an equal reduction in the total contributions of others.

This is all so familiar as to be boring. Sadly, however, although it is boring, the reasoning is also entirely valid. And it is presumably *exactly* what Stern and Garnaut (and economists generally) have in mind when they say that global warming/climate change is the most difficult problem the human race has ever had to face. With all due respect, that problem is *not* the problem as to whether carbon taxes or cap-and-trade or hybrid price/quantity restrictions represent the best mechanism for Australia's policy on this matter. Those latter issues only make it onto the agenda when we have a conceptual solution to the global public goods problem. I say 'conceptual solution' here, because I don't have in mind a requirement that we have an actual treaty with real teeth to which every nation has signed on and to which we have some reason to expect faithful compliance by all players. I mean by 'conceptual solution' that we have at least *some* idea of what would actually motivate the governments of nations to act in ways that are systematically contrary to their *national* interests! In the absence of an answer to

that question, I have to say, it seems like just another case of the can-opener, all over again!

The contrast with *domestic* environmental issues is that these issues are, at least in principle, something we can do something about. We may not know exactly how to measure the marginal benefits of, say, restoring the marginal wetland. We may concede that negotiating the political process (with the complications of our Federal structure) in order to secure what we believe to be genuine environmental gains is a tricky business. But at least there is a political process to negotiate. The policy ends are something Australia can deliver on — and policy-makers can be held politically accountable on their performance in delivering those ends. There is, in short, a collective institution, the Australian polity, that can be an appropriate addressee of any policy recommendation — that can tax citizens to provide the means to buy back water; that can define and assign property rights and enforce the terms of any exchange; that can experiment with the modification of subnational political and social boundaries in the manner that David Brunckhorst recommends.

Generally I don't believe in categorical distinctions: I deal a lot with philosophers and for my taste they use categorical distinctions much too much. I prefer to think of differences in terms of positions along a spectrum. But here, I think, there is something close to a conceptual knife-edge. Given what I take to be a fact, that national boundaries are a prevailing institutional feature of the current world order, I think national and global policy are of different 'institutional kinds'. They are as conceptually separate, perhaps, as the modern state is from Hobbesian anarchy.

On the basis of this distinction, I want to say a little more about both sides of this divide — though, slightly apologetically (for reasons that may already be clear and which I shall restate at the end), mostly on the global side.

12.2 The global aspect

There is, it seems, a consensus within the relevant scientific community: first, that there *is* 'global warming' (or at least 'climate change') of significant magnitude; second, that it is *caused by* increased CO_2 emissions — or, at least, could be substantially ameliorated by a significant reduction in CO_2 emissions; and hence, third, that the global consequences of doing nothing about CO_2 emissions are potentially catastrophic.

Apart from a kind of habitual scepticism about predictions concerning the end of the world, I have only the very weakest grounds for questioning that 'consensus'.[2] It seems to me that at this point we ought take the science as given. And that is what I shall do here.

But I also think we ought to take, as presumptively authoritative, the no less broad consensus within the economics profession concerning incentives to free-ride in public goods provision, that I have already discussed. By way of summarising that consensus, let me put the point starkly: *if the Australian government were an utterly faithful agent of its citizens' interests, it would simply free-ride in the matter of CO_2 emissions.*

In the face of that latter consensus, I confess that I found the conclusions of the Stavins paper puzzling. As he observes, there seems to be some evidence that countries are acting unilaterally in the face of global emissions. Many did sign on to the Kyoto protocol (some, like Australia, more reluctantly than others). Europe is developing its own policy — just as Australia now seems to be. There are good prospects that the United States will do something, despite the resolute opposition to the Kyoto approach. California has its own purely state-based proposal. Broadly, and despite his acceptance of the necessity for government action *within* nations, Stavins seems optimistic about the prospects for decentralised unilateral action *among* nations — with cooperation emerging perhaps in due course. I find this puzzling because Stavins must ultimately be claiming that, while individuals cannot solve n-person prisoner's dilemma problems *qua* individuals, they *can* go much closer to solving them via what we might call the 'partitioning' solution. That is, we partition the set of individuals into two hundred or so groups, such that there is coercive power *within* groups, but where the relation between groups is one of independent action, just like the relation between individuals prior to partition.

That there might be such a solution to the public goods problem strikes me as an interesting speculation, and I shall want to explore it a little in what follows. But I have to say that I know of no formal treatment of any such speculation in the literature.[3] And I cannot find any defence of it in the Stavins paper itself. So we are left with a puzzle. It is one I want to engage briefly.

---

[2] These grounds relate to some niggling doubts about incentives and selection biases in the funding of science.

[3] One might claim that a 250-person prisoner's dilemma is easier to solve than a several billion-person prisoner's dilemma. That I freely acknowledge. But I do not think that that entitles us to think that the 250-person version is solvable. And in any event, the standard solution involving explicit cooperation among potential contributors is not what Stavins has in mind.

I said that if governments were faithful agents of their citizens' aggregate interests, they would not act to curb CO_2 emissions. But virtually no serious scholar of politics I know, whether of the standard 'public choice' school or of the revisionist variant of rational choice political theory that I favour, thinks that government actions faithfully reflect aggregate citizen interest all the time.

The standard public choice account

Take the standard public choice argument first. As a whole slew of regulations and tariff 'protection' measures testify, democratic politics is hospitable to policies that serve to redistribute to well-organised interests away from relatively unorganised ones, at the expense of aggregate interests. Those same general forces that make genuinely free international trade such an elusive policy goal can conceivably be mobilised to promote CO_2 emission reductions. As I see it, orchestrating just such a mobilisation is the agenda that Warwick McKibbin sets for himself in his 'hybrid' approach. Although he declared himself last evening to be dissatisfied with the 'trade' analogy, I think the tariff parallel is useful. Rights to emit, like tariffs, serve to create rents. Who gets those rents is an artefact of the particular regulatory instruments used. McKibbins' aim is to construct those instruments and their allocation so as to establish a more or less stable political coalition supporting the maintenance of the policy that gives those instruments their value. His strategy is rather like that of a company, seeking tariff protection, which gives out shares to a well-devised majority coalition so that that majority will reliably support that tariff in future elections. As our experience in trade negotiations indicates, the political forces supporting 'protective' measures can be extremely stable and impervious to change even when it is in the country's *aggregate* interests to change them. McKibbins' ambition is to do for emissions restriction instruments what jolly Jack McEwen attempted to do for tariffs — create a regime of 'protection all round' in which a critical mass of political forces see themselves as having a stake in the maintenance of that regime. That will serve to embed a policy that will impose net costs on Australia — but it will distribute those costs in a way that is politically profitable.

Put the issue a slightly different way. Just why are Australian business interests generally in favour of the current round of carbon emission entitlements? The answer is surely that they expect some significant proportion of those entitlements to be *given away* to existing businesses. Contrast the creation of these emission entitlements with a general carbon tax — something that Henry Ergas argues persuasively could be a rather better arrangement in efficiency terms, assuming that

most of the revenue were used to substitute for high marginal income tax rates.[4] The carbon tax creates revenues for government. Those revenues can also be given out to special interests — but how that is done remains at the discretion of the government of the day and so is subject to future change. The emission entitlements by contrast effectively assign the revenue value of that carbon tax to those who are given the entitlements: the rent transfers are an intrinsic feature of the policy instrument. Of course, under the carbon tax, exemptions can be granted. But it is by no means easy to disguise tax exemptions; whereas the precise pattern of giving out entitlements is not a salient policy feature. The McKibbin scheme involves, as I understand it, giving some proportion of the entitlements to ordinary citizen-voters, as well as to business interests, presumably with an eye to building around the carbon scheme a robust 'coalition of the willing'.

In any event, this is my reading of the current preference for retradeable permits over taxes. It might all seem somewhat Machiavellian — but another way to read it is as shrewd politics. We create some monopoly-cartel rents in the CO_2 policy process as a means of buying business approval. All this, of course, for a policy that creates net harm for Australia — though if the science is right, probably net benefits for the rest of the world.

If this were all there is to be said about democratic political process, it would not be good news. Taking a step back from the particular environmental application, the logic suggests that it is possible to create political support for almost any policy by a strategic manipulation of the redistributions to which that policy gives rise. If McKibbin can do it for carbon emissions, why couldn't any special interest do it for whatever madcap scheme happens to be on their agenda?

One important part of the answer to this question lies in the constraints imposed by broad public opinion, as revealed in electoral processes. It is one thing to be able to construct a policy framework that will buy off special interests. It is another to do so in a manner of which the electorate will approve. Jack McEwen can deliver 'protection all round' only if there is a general mistrust of free trade and a climate of community support for tariff regimes *in principle*. The standard public choice line is that securing electoral approval does not involve any test independent of the distributive structure of policy, because individual voters can be 'bought off' by strategic redistributions in the way that special interests can. McKibbin's scheme to involve a critical mass of ordinary citizens in the allocation of emission permits suggests that he endorses that standard line. (Note that this involves finding a critical minority of Australians who are going to have to bear the full burden!) But I think that the standard line is at best partial and at worst misleading. Indicating why

[4] Which interestingly in the Australian case does not mean the rates on upper incomes so much as rates on welfare recipients with sharp means tests.

brings me to my revisionist 'expressive' account of voter behaviour; and to a second line of hope for a solution to the global emissions issue.

Expressive voting

The expressive account of voting offers a distinct reason why policies may not be in the aggregate national interest — because votes do not reliably track the voters' individual interests.

The account takes as its point of departure the fact that no individual voter can reasonably expect to be decisive — that the probability of my vote determining who wins in an election is asymptotically negligible. This means, among other things, that voting in accord with my conscience becomes a pretty cheap activity.

Suppose that there's a policy that, if implemented, will cost me $10 000 a year. I think that it's the right policy from a global point of view. But $10 000 a year is a lot of money. On the other hand, what does it really cost me to vote for that policy? Not $10 000 a year. Rather, $10 000 a year times the probability that my vote will determine the outcome of the election! To simplify just a little, $10 000 a year times the probability that there will be an exact tie among all other voters! (In all other cases, my vote is outcome irrelevant: if I made a mistake in voting and somehow voted for my less favoured candidate, that mistake wouldn't actually change the outcome!) Now, the probability of an exact tie among twelve million other voters is a small number. (Actually, the relevant magnitude is the probability of an exact tie in the marginal electorate in an election won by one seat — times the probability that my electorate will turn out to be the marginal one!) This probability is almost certainly small enough to make the cost to me of voting my conscience on this $10 000-a-year matter something like a mere dollar or two. Acting as my conscience dictates may not be worth $10 000 a year; but it is likely to be worth a few dollars.

What this means is that if individuals were truly rational (and held rational beliefs about the probability of being decisive), individual interests would predictably play not much role in politics — and certainly a much smaller role than they do in market settings, where agent choice *is* decisive over options. The right way to think about voting behaviour is in terms of cheering at a football match — not choosing a car or a house or an assets portfolio. When you cheer at a football match, you express your desire that a particular team win: you show your support for your chosen team. But your cheering is not causally efficacious — it exercises negligible effect on the actual outcome. Voting is more like a 'speech act' than it is like a market action — closer to an opinion poll than, say, selecting your portfolio manager!

A direct implication is that the things that *are* relevant in electoral competition, and hence in determining electoral outcomes, are the factors that make individuals cheer (and boo). And one thing that lots of people are likely to cheer for is the environment. Who, after all, wants to express support for global catastrophe?! On this basis, large numbers of voters can *quite rationally* vote for carbon emissions policies that will make them individually considerably worse off (and almost certainly Australia worse off as a nation) because they believe that it is the morally and globally responsible thing to do. And expressing that moral position in a context where they are just expressing their views is much cheaper than doing so in a setting where each only pays if she speaks up.[5]

It is therefore by no means inconceivable that enough citizens in enough of the democratic countries around the world will support policies that inhibit CO_2 emissions entirely unilaterally — and that emissions *will* fall to a level such that potential catastrophe will be averted. And this notwithstanding the fact that the national interest does not support emission reductions in any of the countries involved.

I say that this is not inconceivable. I do not say that it is especially likely. Personally, I am sceptical. My prediction is that, in the medium term, some countries will have made significant sacrifices, but that the level of global CO_2 emissions will nevertheless have increased — with its attendant climatic effects. When that happens, enthusiasm will start to run thin: pre-emptive suicide is not, after all, an especially popular policy. (This, we might observe, is the fate of many popular enthusiasms — like wars that go on longer than a year; or megalomaniacal public projects when the real cost comes home to bite. There is a flurry of electoral support when they are first introduced, but unless there is manifest 'success' the support is hard to maintain.)

Moreover, the danger in expressive politics is that it encourages symbolic policies. I do not, for example, regard it as at all surprising that Kyoto had little overall effect on CO_2 concentrations. Drew Collins made the remark in passing that 'much environmental goodwill had been squandered on tokenism' (nice turn of phrase, that!); but again, that is what I think we ought to expect. For example, when I remarked earlier that, if the Australian government were a totally faithful agent of its citizen's interests, it would simply do nothing about carbon emissions, I did not

5 The line of argument at stake in this view has been elaborated and defended extensively in Brennan, G. and Lomasky, L. 1993, *Democracy and Decision: The Pure Theory of Electoral Preference*, Cambridge University Press, New York. It has implications for policy that are extremely general. Of course, this is not the context to spell those implications out in even minor detail. The remarks here will, though, perhaps be enough to suggest that the expressive account of voting offers a better account of environmental politics than do rival accounts of voting that treat the voter as if she believed that her vote actually determined the electoral outcome!

claim that that is what the Australian government would *say* that it was doing. On the contrary, the demands of international respectability almost certainly require countries to pay lip-service to international environmental agreements; and even to be active promoters of the development of such agreements. And at the domestic level, there will be a lot of emission reduction rhetoric; but there will also be a lot of trying to manage a symbolic commitment to global environmentalism without imposing too great a cost on ordinary producers and consumers. We have already witnessed a little of that here in Australia in terms of the petrol price debate.

In short, I predict a lot of 'I'll fumble; you pay' across the international community; and perhaps even greater than usual amounts of political hypocrisy across the Western world, as self-imposed CO_2 emissions regimes start to bite. That itself presents an interesting policy challenge. 'Hypocrisy' doesn't, I like to think, come naturally to economists — so balancing the requirements of political respectability against considerations of the national interest, truly conceived, is no small ask. And I sympathise with those who are going to have to negotiate the questions. Of course, many here are familiar with that challenge; but I suspect it is going to get worse — much like the environment itself.

From the global to the local

I have spent almost all of my word allocation (and time) talking about the global case. This instantiates what I often say in meetings like this: that there is an important difference between academics and policy makers. We academics can focus on what is interesting (and I have to say that I personally find the global issues fascinating); whereas public servants and policy advisors have to focus on what is relevant. And it should be clear by now that I think the relevant challenge lies with environmental issues that appear at the domestic level.

But although I have drawn a sharp conceptual distinction between global and local, I do not want to suggest that they are unrelated. Indeed, focusing on the connections between the two levels will perhaps serve to sharpen the force of my overall judgement of the proper policy response.

- First, to reiterate the conceptual point. If unilateral action in the global warming case is rational for Australia (or more modestly constitutes a 'political equilibrium' for any individual country), then public goods problems can't be of quite the severity for decentralised institutions that the standard public goods analysis assumes. So, international success or otherwise in the carbon emissions area will have interesting implications of a general theoretical kind for how we think about public goods and common property problems more generally — and specifically at the more domestic level.

- Experience of success or failure in more national/local environment problems can inform us as to the prospects of success in international ones. If the fact that political jurisdictions are overlapping creates problems for, say, resolution of the Murray–Darling problem, when the number of governmental players is small, what are the prospects for tolerable solutions to the carbon emissions issue when the number of players is huge and highly heterogeneous, and there is no overarching political authority?
- The severity of purely local environmental problems is likely to be a function of the more global situation. For example, climate change will affect rainfall levels, with implications for water management in river systems (a prospect that both Arlene Buchan and Wendy Craik note in their papers in relation to the Murray–Darling). More generally, as Gary Libecap notes, norms and policy régimes that have governed management of common property resources reasonably satisfactorily can become deeply problematic when external (climatic) circumstances alter significantly. Arrangements that have been in place for some time acquire the status of de facto 'rights' with all the moral freight that rights carry and all the potential for moral outrage when the rules of the game are changed. (Such moral outrage problems may be exacerbated if the Australian government gives voters an expectation that their sacrifices will actually solve the global problem.)
- Resources deployed by the Australian government for tackling global warming are resources that could otherwise be used for tackling national environmental problems (such as water buy back). I want to emphasise this 'resource competition' both because I think it is empirically very important and because it is an aspect of things that economists are distinctive in recognising. Environmental scientists tend to think of environmental policy as a seamless web of interconnected relations. As Suzi Kerr put it in her presentation, methane gas emissions and nutrient leakage into lake water are similar because they involve the same source of pollution and the same 'science'. But, as I have argued, they also have quite different status in policy terms. And there is a real choice about where we focus our efforts.
- These 'resources' include, at the most general level, not just fiscal dollars (or GDP) but also the scarce attention of politicians and bureaucrats and for that matter, of citizen-voters. For policy purposes, this may be a yet more significant source of conflict.

Ultimately, my advice to the policy advisors, somewhat tentatively given as always, tracks my advice to students in relation to exam strategy: 'start with the easy problems, the ones you think you can solve. Worry about the harder ones later'.[6]

I do not of course pretend that domestic environmental problems are simple. But they are ultimately something we can actually do something about. And that seems to me to distinguish them from their global analogues. Domestic environmental problems are where I see the expected policy payoff to lie. For that reason, I think that the papers presented here today, focusing as they have on the domestic cases — at both practical and theoretical levels — have got the balance right.

Australian policy should focus on domestic environmental issues — and on the global front, focus on *optimal response to the fact of global warming* rather than on Quixotic attempts at prevention. As David Parnell put it, the main game is adaptation, not mitigation. This is not because it is impossible to imagine that the world might luck into a kind of solution to the global emissions problem. I have tried to give reasons why that is not an inconceivable outcome. But I think it distinctly against the odds.

6 Of course, the exam advice comes in a setting where 'all questions are of equal value'. That may not be the case in the environmental area. But I still think that it's a fair rule of thumb not to spend all your effort on problems you are very unlikely to be able to do anything about!