STUDIES IN REGULATION

Re-thinking climate change policies: A tale of two externalities

Gerard Fox and George Yarrow

NS 12.1 January 2022

REGULATORY POLICY INSTITUTE

Studies in Regulation, New Series, Vol 12, No. 1

Published by the Regulatory Policy Institute
300 Banbury Road, Oxford OX2 7ED, UK
www.rpieurope.org

First published January 2022

© Regulatory Policy Institute

Gerard Fox
George Yarrow

Re-thinking climate change policies: A tale of two externalities

Introduction

The purpose of this short paper is to raise a rather fundamental, but neglected, issue of conceptual framing in assessing the conduct of climate change strategy. It is not concerned with issues of science, such as the rate of warming or the links between that rate and concentrations of atmospheric greenhouse gases, nor with the fine detail of policy trade-offs. The level of abstraction is deliberately high, although sufficient to give some pointers for policy alternatives to a current approach that we believe is flawed. Whilst there is both a raft of technical economics behind the points to be made and some obvious developments required for a more rigorous exposition, these have been short-circuited to better to focus on conceptual framing, on how the relevant issues are conceived/pictured in the mind.

The note starts with a stylized characterisation of the thinking lying behind the current fashion for policies that rely on a particular view concerning *economic externalities* (alternatively *external effects*), i.e. those effects of activities on others, for good or ill, that are distinguished by being uncompensated (unrewarded or unpunished). It then examines the difficulties introduced by the current conceptual framing of the climate change issues, before showing that *the nature and significance of any potentially problematic external effects is a function of strategic policy choices*. To connect the abstraction to an observable analogy, we illustrate the alternative conceptualisation by reference to the development of the Oxford/AstraZeneca Covid vaccine. This is followed by some conjectures concerning science and technology policy and a few concluding remarks of a more wide-ranging nature.

The sentiment could be said to be Wittgenstein's, "To show the fly the way out of the fly-bottle" (Philosophical Investigations), i.e. to address self-created obstacles to achieving that which is desired. In the context of relevance here (i.e. assessment of externalities) this is an exercise that was performed by Ronald Coase in what is claimed to be the most cited economics paper of the twentieth century, *The Problem of Social Cost*. The paper also has a strong claim to being the most misunderstood paper of the period.

The conventional quantity setting approach/'wisdom' in a nutshell

We start with the propositions that (i) global warming is, at least past a certain point, harmful and (ii) the harms caused by it are an increasing function of the amounts of GHGs in the atmosphere, usually expressed in terms of CO_2 equivalents. We take no view on the rate of warming, other than presuming that it is positive.

The quantity of CO_2 in the atmosphere is a what is called a state variable (or, in older economics terminology, a stock variable): it has no time dimension to it, unlike say a MWh of energy, one megawatt *per hour*. As such, the quantity of atmospheric CO_2 at any point in time will be determined by a cumulation of *net* emissions over earlier periods.

Within this framework, net emissions over a defined period, usually a year, are cast in the role of control variables and they become the centre of cognitive attention: control net emissions and that will determine the future course of atmospheric concentrations of CO_2. Other things equal, reductions in net emissions will, cumulatively, reduce future levels of atmospheric CO_2; increases in net emissions will increase those future levels

Next comes an Economics 101 conceptualisation of the issues. Emissions levels in any year are typically evaluated relative to some prior, benchmark levels, e.g., levels recorded in 1990. Reductions below this level can be referred to as 'abatement' and the basic proposition is that there are both benefits and costs associated with abatement. The benefits stem from avoidance of future harms that would be caused by otherwise higher atmospheric CO_2. The costs come from the technological and behavioural adjustments required to increase abatement, i.e. to reduce emissions. These benefits and costs of higher abatement can therefore be traded off against each other in determining an abatement level that would maximise net benefits.

It should go without saying that the assessment of benefits and costs is fraught with uncertainties and the idea that any 'optimal' path could be identified with precision is utopian. However, the stylized notion of a balancing judgment that necessarily has to be made suffices for the purposes of this current exercise.

Classroom exposition typically examines the abatement level that would give rise to the highest net *global* benefit in an upcoming period, $B_g(X) – C_g(X)$, where X is the global level of abatement and subscript g indicates that the relevant benefits and costs are to be calculated at the global level. The condition for doing this is that marginal benefit, MB_g, the increment in benefit from a small increase in abatement, is equal to marginal cost, $MC_g(X)$, the extra cost of achieving that small increase.

In this conceptualisation of the policy challenge the problem arises in consequence of the harm-causing factor being global in nature whereas policy actions are determined by national governments. Each government will look at the problem and note that, if it reduces emissions by some target amount, it itself will enjoy only a fraction of the benefits that will accrue from that action. In contrast, acting unilaterally, it will bear all the costs of the action.

Thus, a nation the size of the UK could reckon that it would enjoy only around 1% or less of the additional benefits resulting from its own efforts and that the other 99+% would be enjoyed by the citizens of other jurisdictions. The disparity would be less for larger emitters, but would still be only fractional. We will call this 'Externality Q' (Q for quantity of abatement) and the expectation is that abatement levels would be seriously sub-optimal if nations proceeded unilaterally on this basis.

An expectation of serious sub-optimality also applies in respect of any international agreements that might be made and which have the characteristic of being 'self-enforcing', meaning that once made there are no (or in practice highly muted) incentives for any party to deviate from the agreement. In economics and maths such agreement points are referred to as 'non-cooperative' Nash equilibria, although 'non-cooperative' is a slightly misleading term since an agreement might be needed to achieve them. 'Limited co-operation Nash equilibria' would arguably be a more appropriate label.

Such equilibria were much studied by military strategists in the Cold War era when the underlying issues at stake concerned trust, incentives for compliance with any agreements made, and verification (by of their nature, these equilibria addressed the compliance problem). Whilst more ambitious agreements might be feasible, if, once that is done, *pacta sunt servanda* is an unreliable basis on which to proceed, other incentives for compliance (e.g. punishment mechanisms) are relatively weak and verification is challenging, any agreement struck (on arms reduction) that went well beyond the Nash equilibrium could be expected to be fragile.

An additional difficulty: strategic substitutes in quantity-setting games

Economists might note the similarity of emissions targeting framework to that of the Cournot model of duopoly or oligopoly to which undergraduates are routinely introduced. In both cases the decision agents determine the *quantities* of something. For Cournot businesses the something is the output they plan to produce and offer to the market, for national governments in the climate change context it is the level of national abatement.

In graphical expositions of the Cournot model, the outputs are typically assumed to be what have subsequently been characterised as *strategic substitutes.* That means that if, for whatever reason, one firm increases (decreases) its output, the optimal response for a rival is to reduce (increase) its own output. It is not a necessary implication of the quantity-setting approach: as always things depend on the shapes of demand and cost curves, but it is usually a safe assumption for empirically normal curves. Economists familiar with the model may remember that graphical illustrations in the textbooks of the response functions (alternatively reaction functions) show them as having negative slopes: if a rival does more (sets a higher output), a responsive business will do less (set a lower output).

In the economics canon, the Cournot model has been extended to the Stackelberg case, in which one business (the leader) makes a pre-commitment to a given level of output and others follow, taking the leader's commitment as a given. The taken as given assumption defines a *residual demand curve* for the followers, a concept that perhaps enjoys most systematic, current usage in the analysis of the business conduct of suppliers in wholesale electricity markets, particularly when analysing an electricity system operating in stressed conditions, i.e. conditions in which total demand is close to maximum available generation capacity.

In a context of quantity-setting (setting of abatement levels) in climate change policies, the Stackelberg case corresponds to circumstances in which a single country self-identifies as a leader and makes credible commitments to its own abatement strategy. In these circumstances the relevant decision on the quantity to be achieved is not contingent on any reciprocal actions by others (unlike in the negotiated settlement interpretation of the process that leads to the Cournot/Nash equilibrium). To a reasonable approximation, successive British governments have seen themselves as such a leader, albeit without much due diligence in exploring the consequences of doing that.

Whatever the motivations for such a move – whether a sense of moral purpose, or virtue signalling, or *pour encourage les autres* – the effect is the same and is shown in Figure 1 below. The relevant, *national* ('residual') marginal benefit curve (MB_i) is that of a 'follower' nation that determines its own abatement level x_i (lowercase x signifying a national abatement level and subscript i identifying the individual country). Similarly, MC_i is the marginal abatement cost curve for nation i.

The marginal benefit curve is shown as relatively flat, as would be expected for a country that accounted for only a modest fraction of global emissions, reflecting the fact that its own emissions make only a small contribution to the global total. The intersection of the curves/lines determines the abatement level that maximises *national* net benefits, initially point A. The effect of the leader's commitment to a higher level of abatement than previously planned is to shift this curve downwards, leading to an adjustment to point B, where the follower's abatement level is lower than at point A.

Figure 1, Follower's quantity-setting reaction

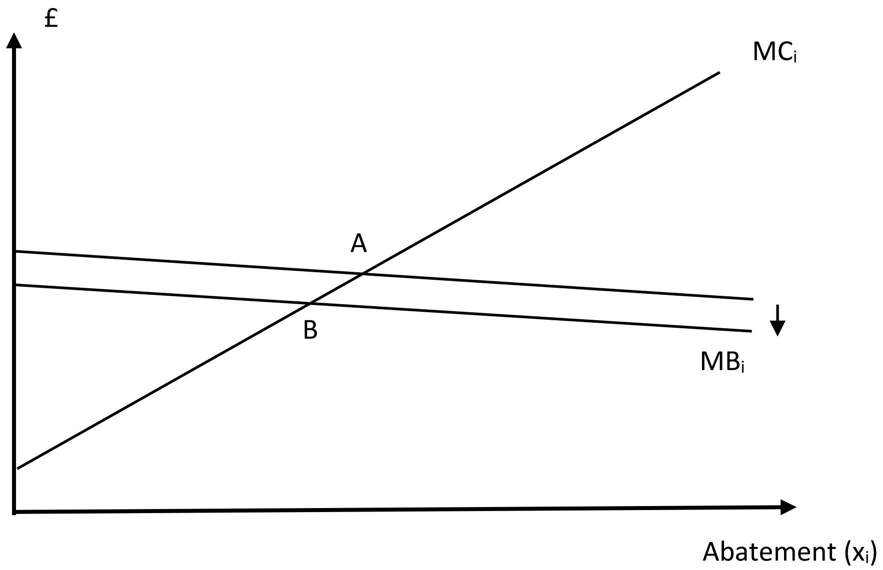

A displacement or leakage effect can therefore be expected to eventuate, and this is what is meant by saying that quantities are strategic substitutes. To achieve, say, an extra 1000 tonnes reduction of atmospheric CO_2 a 'leader' will need to reduce its own, national emissions level by something more than that. If, for example, the responsive actions of others amount to reductions in emissions equal to 25% of the leader's incremental abatement, the leader would require a 1333 tonnes reduction in its own emissions to achieve the 1000 tonne global reduction sought. With rising marginal costs of abatement, leadership could therefore prove a very expensive strategy to pursue.

An alternative perspective: the economics of information

Governments cannot in general simply mandate levels of national abatement and, if seeking to achieve quantitative targets, they tend to have resort to a wide variety of interventionist measures aimed at inducing the desired effects. One of the most prominent areas for such interventions is science and technology policy, particularly the promotion of cost-reducing innovation in low-carbon technologies, ideally to a point where these technologies become lower cost than carbon intensive alternatives. Since costs are measured in currency units (£ per unit, $ per unit), such policy is addressed at changing the *value/price* of something, not its quantity.

To the extent that they are successful, such policies will shift the marginal abatement cost curve downwards, albeit likely by incurring fixed costs arising from the initial development expenditures. This is shown in Figure 2, which illustrates the position of a technology '*leader*' or, more generally, a nation that chooses to increase (above some prevailing level) resources devoted to abatement-cost reduction on a *unilateral* basis.

Other things equal, the effect of a downward shift in the marginal cost curve is to lead a net-benefit maximising government to *increase* its level of CO_2 abatement, so the policy can be viewed as an *indirect* means of achieving the aspiration of reducing emissions. This is a classic strategy point: the introduction of a strategic objective (in this case reducing the costs of deploying low-carbon technologies) which, if achieved, contributes positively to a wider, overarching aim or aspiration. In a climate change policy context, this particular strategic objective is interesting for some rather fundamental reasons, which have to do with the economics of information.

Figure 2: Technology 'leader's' position

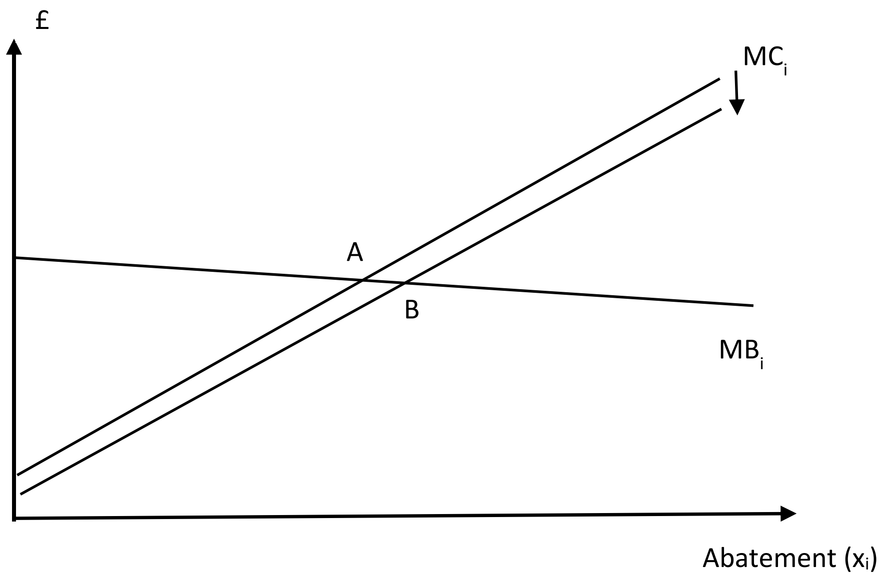

Viewed as a commodity, Information/knowledge has rather specific characteristics. One is that production of it is by way of a *discovery process*, the outcomes of which are, definitionally, unknown. It could almost be a song lyric: "you don't know what you'll get 'til it's found".

This is uncertainty in the raw and it can pose severe psychological challenges for those who would want to see a relatively high degree of certainty about outcomes, e.g. achieving pre-defined targets developed in some or other planning process. The approach (promoting abatement cost reduction) cannot be regarded as a simple 'instrument' to achieve what is desired: the target may be underachieved, or it might be overachieved. Rather it relies on the setting of a strategic/indirect objective on the basis that all that can be known is that success in its pursuit can reasonably be expected to contribute to success in achieving an overarching aim.

A second aspect of information/knowledge is that, depending upon its precise nature, it can share one of the characteristics of a 'public good', *non-rivalry in consumption/use*. One person's use of the information does not subtract from its availability for use by other persons. Having been discovered, Newton's laws of motion are information/knowledge available to anyone, and anyone can read the script of Hamlet.

Unlike a pure public good, others can potentially be excluded from its use – it can be kept secret, for example – but there are significant, sometimes very substantial, costs of doing that and, in practice, valuable new information/knowledge tends to spread via processes of non-transactional diffusion. The discoverer of the new information, who may have put considerable effort/resources into its discovery is financially uncompensated for these (non-transactional) benefits that it will provide to other users of it. We will call this ‘Externality V’ (V for value).

Considered in isolation, Externality V works in the same direction as Externality Q (the uncompensated beneficial effects conferred on other nations by increased abatement in any one nation). If the activity (whether carbon emissions abatement or discovery of new information/knowledge) is financially costly, there will be incentives to supply a sub-optimal quantity of it.

Things look very different, however, if the two externalities are examined in conjunction with one another. To see this, return to the nation-level expression for marginal benefit as a function of carbon abatement and assume, for simplicity and as illustrated in Figures 1 and 2, that the national curve of country i (MB_i) takes a linear form:

$$MB_i = \theta\, MB_g(X) = \theta\, [\, a - b\, (x_1 + x_2 + \dots + x_n)\,]$$

where a and b are positive parameters that together define the location and position of the curve and θ is the fraction of the global net benefits enjoyed by nation i (which might take a value of around 0.01 or lower for a country the size of the UK).

Now suppose that a country puts greater effort/resources into promoting the discovery of technologies that will lower abatement costs and that it achieves at least some success with that policy. If the advances in knowledge are made available to other nations, each of the others can make valuable use of them. Their marginal abatement cost curves will, like the leader’s shift downwards.

Mathematically then, we can, because of Externality V and with other things equal, expect the abatement level of each country to be an *increasing* function of the discovery efforts of any single country i, denoted y_i. That is:

$$MB_i = \theta\, MB_g(X) = \theta\, [\, a - b\, (x_1(y_i) + x_2(y_i) + \dots + x_n(y_i))\,]$$

Thus, when country i increases its discovery efforts and achieves successes from that, the effect is to increase abatement activity *everywhere*. Externality Q is now reversed: country i, which may or may not be a self-identifying ‘leader’, benefits from the induced, incremental abatement that takes place elsewhere in the world as a result of the availability of the new discoveries. A country that makes an extra effort in discovery is, if it is successful, thereby compensated for the effort by the consequential actions induced in other countries. It is not financial compensation, but rather a reward ‘in kind’, taking the form of increased abatement levels in other countries. *Precisely because of the existence of Externality Q*, these will confer benefits on country i.

The underlying point is a very simple one: if, by seeking to reduce emissions, any one country wishes to protect itself against possible harms from climate change caused by the build-up of CO_2 concentrations in the atmosphere, *it really doesn’t matter where in the world the reductions in emissions are made.*

The question that can be put, then, is: *Where’s the externality?* It may be, of course, that the compensation enjoyed does not fully compensate for the costs of the extra discovery efforts made unilaterally by an individual country, but then again it could possibly be the case that those efforts are over-compensated, i.e., the value of the benefits received may significantly exceed the costs that have been incurred. It all depends upon the specific features of the relevant, factual context: a major scientific breakthrough could, for example, easily have an overcompensating effect. What cannot be substantiated, however, at least without much more detailed empirical assessment, is that there exists a net externality that *necessarily* gives rise to globally sub-optimal emissions abatement in the absence of a global agreement.

Strategic complements

The policy implications of these points are potentially major. If policy strategy is focused on making commitments to abatement levels (*quantities* of emissions), upping the commitments can be expected, other things being equal, to have at least some 'crowding out' effects on the contributions of others (the *strategic substitutes* point). On the other hand, if the policy strategy is focused on cost-reducing innovations, the non-rivalrous nature of information/knowledge tends to lead to *strategic complementarity*. Successful, innovative efforts in cost reduction by one nation will 'crowd in' abatement cost reductions, and hence higher levels of abatement, in other nations: their cost of abatement curves will be shifted downwards.

This effect will be reinforced if, in addition, there is strategic complementarity in *discovery efforts* themselves, i.e. one nation increasing its effort will lead others to *increase* their own efforts. Determining whether or not that is likely to be the case depends upon finer-detail contextual analysis and is therefore left as a question beyond the scope of this paper.

In respect of these matters, there is an analogy with theories of duopoly/oligopoly taught in undergraduate economics. The Cournot analysis was first published in 1838, but in 1883 a critique of it appeared in a book review by Joseph Bertrand. Bertrand's model switched the analysis (of the same business context) to a focus on the choices of the prices (values), not the quantities, of the products that the businesses produced and supplied. That is, it examined the non-cooperative Nash equilibrium of a price-setting game, rather than a quantity-setting game. The results that followed were strikingly different.

When (a) business strategies are cast in terms of price-setting and (b) the relevant products are differentiated to at least some extent, the response/reaction functions tend to slope *upwards*. If one firm takes the lead in committing to reduce its price, other firms tend to follow it down. If the commitment is to a higher price, other firms tend to follow it up.

In each of the two models all players are assumed to compete in a similar way – either all set quantities or all set prices – but in actual business contexts firms can usually choose *how to compete*, and later technical work in economics has explored some of the implications of the existence of a quantities/prices choice. This is a higher level 'game' and it brings abstract game-theoretic analysis a step nearer to business strategy analysis.

The most important propositions of this short paper in the climate change context are, therefore, that similar (quantities-or-values) strategic choices are available to policymakers and that the choice actually made can be expected to matter a lot. It can also be added that, in practice, the climate policy choices available are not of a simple binary nature: governments can, and do, pursue both quantity-reducing and cost-reducing objectives simultaneously. The challenge then becomes how to find/discover coherence and balance in the mix to be chosen.

Considerations of strategic substitutability and complementarity are highly relevant to assessment of the policy balance – as argued, choices made affect the headwinds and tailwinds (the various external effects) – but so too are issues of policy coherence. In complex economic contexts where a range of regulatory interventions are in play it is usual to find major problems of what can be called *regulatory* *interference.* Roughly, policy X may be aimed at achieving outcomes falling in a set A and policy Y at outcomes falling in a set B, but each policy may impede the effectiveness of the other. The interference problem is ubiquitous in public policy and, for current purposes, we simply note its relevance in the climate change context and pass quickly on. (It does, though, define a rather comprehensive assessment agenda for any government interested in raising its own regulatory performance.)

The Covid vaccines analogy

To illustrate how the above, abstract points might quickly translate into strategic policy thinking, consider the development of the Oxford AstraZeneca Covid vaccine, intentionally developed as a *vaccine for the world*: payments for Intellectual Property Rights (IPRs) were waived for an initial period.

The primary motivation for the strategy in this case was a moral and cultural one, very much in line with the kind of ethos to be found in the university environment in which those who developed the vaccine worked. A university is (or should be) an institutional form geared toward the discovery *and transmission/sharing* of new knowledge. However, in later justifications of the policy/strategy a version of the self-interest (where's the externality?) point set out above did appear.

Thus, it was argued that the defeat of the SARS-Cov-2 virus in countries other than Britain was itself beneficial for the people of Britain: the virus is highly transmissible and hence, for as long as it was circulating at a high level among citizens of other countries, any contacts between them and British Citizens risked the transmission of the virus from elsewhere to Britain. Risks of new variants emerging have also been mentioned in support of this self-interest justification for promoting wide availability. In limiting the spread of the virus elsewhere, therefore, there were at least some 'feedback' benefits to the UK.

The feedback effect can, however, be expected to be substantially weaker in the vaccine context than it is in the carbon abatement context. In the latter case it doesn't matter at all where the induced effects overseas are located. A given level of emissions reduction would have the same benefits for UK citizens if it was made in Brazil, or South Africa, or India, or Australia, or North Korea, or wherever. This is not the case in relation to the Oxford AstraZeneca vaccine, where, for the purposes of the relevant self-interest argument, geographical differentiation is a factor that clearly needs to be considered.

That is because the rate of transmission of the virus is a function of the frequencies and durations of close social contacts and these vary considerably, both within and between nations. It is, for example, more likely to be beneficial for the welfare of Britons to vaccinate a British citizen than, to take an extreme comparator, a citizen of North Korea. Nevertheless, the basic, 'feedback' argument does still apply in the Covid context: it is simply a matter of the chains of causality being weaker and much more differentiated/variable.

The vaccine analogy also highlights another potential factor that has bearing on climate change policy in that, in competition with the Oxford AstraZeneca vaccine, pharmaceutical and biotech companies motivated by financial returns stepped into the ring with their own products, some of which have found wide international markets. This contemporaneous development points towards a number of interesting issues for science and technology policy, as we now discuss.

Trade-offs in science and technology policy: the transmissibility of information/knowledge

Returning to the central theme, there is, we think, a rich agenda for policy thinking in further, extensive development of the information/knowledge sharing points that have been made. As is the case when thinking about the implications of the distinction between strategic substitutes and strategic complements, a basic economics can provide helpful, foundational material, but it tends not to go deeply enough into the specific detail of individual policy contexts.

Thus, the more basic theorising about 'non-rivalry' in the use of information tends not to get to grips with the observations that (a) there are different types of information/knowledge and (b) these are not equally *transmissible.* Newton's Laws of Motion are highly transmissible, but the kinds of information that can be called 'know-how' are much less so. Indeed, the non- or limited-transmissibility of the latter plays a key role in the dynamics of competition by providing a source of competitive advantage to those who can acquire it. The competitive advantage arises from precisely the fact that it is not easy for others to acquire and make use of it.

As well as inadequate depth of thought in the theorising, there tends also to be a lack of breadth. As mentioned above, consideration of externalities tends to be partitioned into boxes, which is understandable in the teaching rooms of universities where the focus may be on developing analytical techniques. Yet policy challenges don't present in unbundled ways. In practical contexts, externalities tend to *"come not as single spies, but in battalions"*. An adequately responsive analysis therefore needs to be 'holistic' in its approach, by which we mean that all salient features and behaviour patterns in the relevant context need to be taken into account and all salient interactions need to be explored.

Perhaps nowhere is the muddle caused by neglect of these points more currently visible than in the conduct of science/technology and industrial policies in the context of their new relationships with environmental policies. Questions that might be asked about any new programme include: 'What precisely is the aim here? Is it to reduce global carbon emissions? Is it to achieve competitive/comparative advantage for the country? Is it to 'level up' within the country?

A politician's answer is likely to be that multiple aims are encompassed, in which case a variety of other questions arise covering matters that include assessments of trade-offs, priorities, appropriate strategic objectives, the various different sets of mutually coherent actions that might be taken, institutional frameworks, and so on. If reducing global carbon emissions is the major priority, that points to promoting the discovery and free sharing of highly transmissible information/knowledge. If gaining competitive/comparative advantage is also a priority, that points to promoting the discovery of much less transmissible knowledge or of not freely sharing transmissible knowledge, but rather charging for it. It is easy to see, therefore, that the different directional pointers implied by multiple objectives can lead to confusion and to regulatory interference (as defined above).

In relation to the institutional issues, for example, a government that wanted to pursue both aims simultaneously might usefully think about the option of having recourse to a division of labour/responsibilities within government based on the transmissibility of different types of new information/knowledge that are expected to flow from different discovery processes in differing contexts. A Climate Change Agency might, for example, be afforded responsibility for those aspects of science and technology policy that were focused on transmissible information/knowledge relevant for carbon abatement, whilst a Business Department might have a focus on the more traditional tasks surrounding the development of competitive/comparative advantage.

It is to be recognised too that transmissibility is something that science/technology policies can influence: it is not necessarily an inherent characteristic of the new information/knowledge itself. The vaccine analogy is a good example of that as the different marketing strategies adopted by AstraZeneca, Pfizer and Moderna illustrate. The time-limited waiving of IPRs for the Oxford AstraZeneca vaccine was very much a matter of choice, a decision to promote a greater uptake elsewhere in the world.

At the institutional level, the division of responsibilities/tasks within the system of government departments and agencies is another matter of strategic choice, best made with a view to establishing clear mappings between responsibilities/tasks and factors such as relevant purposes/priorities and the likely transmissibility of the type of new knowledge sought. Strategic adjustments called for by changing circumstances could then be reflected, in a traditional way, in variations of the budgets of different departments and agencies.

As in relation to the more general regulatory interference problem, there is another potentially very large agenda to be addressed in these areas.

Concluding remarks

The idea that public policy should take account of an existing, 'given' (i.e. not avoidable) externality when considering actions to address another, identified externality is not a new one: in economics it goes under the name 'second best theory'. However, that label puts it into a sort of cognitive box where it becomes something of a side issue, whilst vaguely insinuating that 'first-best' is a feasible goal and is perhaps a more interesting area of analysis, both of which it is not. What is identified as 'first-best' is typically derived from a modelling ideal rather than an assessment of practical realities.

More generally, the policy analysis of externalities has tended to be used for relatively narrowly focused issues, taking account of one troublesome problem at a time and not infrequently accompanied by a declaration that a 'market failure' has been identified and that it needs some government attention. Indeed, Sir Nick Stern, lead author of the Stern Report in the UK, said (in 2007) that *"Climate change is a result of the greatest market failure the world has seen. ... The problem is global, and the response must be a collaboration on a global scale."*

Each of these two statements is, we think, challengeable. The first, at least with anything like the force with which it was put, begs questions that can only be settled by a more holistic examination of external effects. In respect of the second, in our own view, the top-down global approach adopted has been more of a barrier to progress than a facilitator of progress. It has been yet another attempt to address complex issues via what is intended to be a form of central planning, based on a top-down allocation of quantitative targets. The historical evidence on the performance of this way of doing things is, to put it mildly, not favourable. It is, as it was once put in a paper that was also first put into the public domain in 2007, to wear 'the wrong trousers'.[1]

[1] https://sciencepolicy.colorado.edu/students/envs_5720/prins_2008.pdf

The now decades long distraction from consideration of alternative, likely more effective approaches to the climate change challenges is another reason to question the degree of priority currently given to the pursuit of global agreements that seek an international allocation of emissions targets – and they are only targets, with no reliable guarantees they will be met, or even approximated. This is not to say that they have no role to play, only that they have major limitations and that jaw, jaw (or blah, blah, blah as it has recently been put) is not in itself a sufficient basis for the development of good policy.

Good strategy development looks for points of leverage, where maximum influence or effect can be achieved. They are, we think, more likely to be found in the UK's science and technology base than in global conclaves. It also looks for complementarities, so that activities in pursuit of the disaggregated objectives of individual units of an organisation or system of organisations are mutually supporting, rather than mutually inhibiting. It is difficult to overemphasise the significance of this point. It is key to issues of strategy choice and policy system design.

By way of illustrating one, specific obstacle to the development of effective climate change policies, we draw attention to the accounting system that has been established by the global agreement process (and we simply note it here rather than attempting any systematic examination of its distorting effects).

It is an accounting system based on top-down quantitative planning, taken from the Soviet songbook, which ignores the *value* of important activities that contribute to reductions in atmospheric carbon (just as Soviet-type planning ignored prices). There is, for example, no provision to account for the effects on emissions abatement of cost-reducing breakthroughs made in one country which can be adopted and used by other countries. It is as if Externality V does not exist, but it manifestly does.

'Net Zero in one country by 2050' allocates no credits to nations who manage to achieve and share such advances in information/knowledge. They are unpriced/unvalued in the agreed accounting system and, given the definition of externality with which we started (uncompensated effects on others), there is a great irony in this. Then again, it is not entirely surprising: even casual study of public policy making is enough to indicate that the State is a great creator and upholder of inefficient systems of externalities, even as it may be claiming to be a mitigator of such things (as when it claims to be seeking to correct 'market failures').

To sum up, in respect of climate change policy what we advocate is some cool, clear thinking about what is being done in the name of public policy, rather than having resort to casual moralising of dubious merit or succumbing to a socially viral infection of the messiah complex. Yes, our ancestors have endowed us with a technological inheritance that is now giving rise to serious economic, social and political challenges, but today's inhabitants of earth have also inherited from them a hugely enhanced information/knowledge base from which those challenges can be tackled. Yes, many of the now problematic technologies were first developed and applied in the West, but so also have been technologies (such as, most recently, the computer and the mobile phone) that have proved of great value everywhere.

Large scale, international transfers of economic value deriving from knowledge/know-how have occurred in the past and it would appear to be eminently feasible to sustain a continuation of that process in the climate change policy context. Balanced judgments must rest on a holistic view of *all* relevant, underlying externalities and such a view is not currently to the fore in climate policy assessments.