Comments on “Carbon Valuation in UK Policy Appraisal: A Revised Approach” (DECC, March 2008)

Dr. Frank Ackerman
Stockholm Environment Institute-US Center
Tufts University

April 10, 2009

Summary: These comments have the following major sections.

1. I support the general approach taken to the cost of carbon; the arguments for primarily reliance on marginal abatement costs are persuasive, and stated well in the draft. Estimates of the social cost of carbon are sufficiently problematical and incomplete that they should play little or no role in policy appraisal. The determination of overall targets for greenhouse gas reduction should be, and often is in practice, based on a precautionary response to catastrophic worst-case risks, not calculations of the social cost of carbon.

2. While marginal abatement costs are more straightforward and well-defined, they are still subject to some ambiguity of definition, as recognized in the multiple scenarios in the draft. I suggest a focus on the effects of oil prices and discount rates, which may be more important than some of the variables considered in the current draft. On the other hand, the uncertainty about non-CO2 greenhouse gas abatement, in agriculture and waste, does not seem warranted; these are well-understood areas offering low-cost abatement opportunities.

3. The gap between the price in the traded and untraded sectors seems logical in theory, but is large enough in practice to raise questions about the potential for additional, comparatively low-cost reduction opportunities in the traded sector. If the legal constraints and the problem of leakage to other EU nations can be dealt with, “overcompliance” with the ETC cap seems much more cost-effective than other measures that will be supported by the proposed carbon price in the untraded sector.

4. My comments end with a list of my responses to the numbered questions asked in the specifications for the peer review; the substantive responses are also covered in the body of my comments.

1. Beyond the social cost of carbon

The most important innovation in the draft approach to carbon valuation is the move away from the social cost of carbon (SCC), toward valuation based on marginal abatement costs (MAC). Having advocated the use of MAC estimates for carbon pricing, in comments on past policies, I am a strong supporter of this move. The point is not that MAC calculations are perfectly transparent and unambiguous; problems with MAC estimates are discussed in the next section of these comments. The problems with SCC calculations, however, are insurmountable. As a result, even the limited uses of SCC proposed in the draft are problematical.

The SCC purports to be a single number representing the monetary value of the long-term damages attributable to emissions of an additional tonne of CO2-equivalent greenhouse gases. There are at least three fundamental flaws in the attempt to calculate this cost:

- the precise nature and extent of long-term damages is uncertain;
- the most important damages do not have meaningful monetary prices;
- the long time horizons make the calculation hyper-sensitive to the choice of the discount rate.

Climate change is known to cause many types of damages to economic activity, human health, and the natural environment. The exact timing and extent of these damages, however, is not known and is probably unknowable. The earth's climate is a strongly non-linear system, with uncertain but potentially severe threshold effects, and the possibility of chaotic or complex dynamics – which would defeat any attempt at detailed forecasting. The social consequences of physical climate impacts, such as political conflict and mass migration, are even less predictable, but likely to be very costly.

There is no standard methodology for creating comprehensive estimates of these damages. In practice, published estimates of the SCC are based on scenario modeling or case studies focusing on a small number of well-defined impact categories, such as agricultural losses, property at risk from sea-level rise, storm damages, etc. Studies differ in the categories they include, and in the assumptions they make about those categories; some studies, even among those published recently, rely on outdated primary research about damage costs in particular sectors. The natural tendency is for SCC estimates to be underestimates, due to their limited coverage of the broad range of climate impacts.

One important example is the economic modeling in the Stern Review. Although criticized by some economists for allegedly overstating climate damages, the Stern Review more likely understates the severity of the problem, and hence underestimates the SCC. In an article I co-authored with Chris Hope (the developer of the PAGE2002 model used in the Stern Review) and colleagues, we demonstrate that Stern's analysis using PAGE2002 includes implausibly small damages for the United States and other developed countries. This occurs in part because the model assumes that very extensive adaptation will occur, at low cost; Stern's damage estimates are net of this hypothetical, successful adaptation effort.[1] Thus it is necessary to clarify the adaptation assumptions used in the development of SCC estimates, along with many other definitional issues.

Not only is it effectively impossible to offer a complete description of the anticipated climate damages; it is also impossible to put meaningful price tags on the most important damages. Climate change threatens to cause loss of human life, extinction of species, and destruction of unique, irreplaceable places and ecosystems – none of which are commodities sold in markets. The attempt to assign a price to human life is ethically troubling, as is the subsidiary question of whether to use the same price in rich and poor countries. Yet any comprehensive estimate of the SCC must include a value for the potential loss of human life. I have written at length elsewhere about the ethical and logical problems of "pricing the priceless."[2]

Finally, climate damages resulting from today's carbon emissions will stretch over centuries to come. Quite aside from the uncertainty surrounding multi-century predictions, there is the stubborn problem of discounting: how much are future centuries' damages worth to us, compared to the same damages today? (The familiar but inescapable quandaries in this area are summarized well in Chapter 2 of the Stern Review.) The choice of a discount rate has a remarkable effect on the SCC. In Richard Tol's most recent literature review of SCC estimates, almost half of the variance in SCC is explained by discount rates alone.[3] An increase in the discount rate of one percentage point multiplies the estimated SCC, on average, by a factor of 0.6.

For all of these reasons, the SCC does not provide a stable, reliable basis for decision-making. The limited role proposed for the SCC in the draft document, in setting overall targets for reduction, is no more defensible than its detailed use in policy appraisal in the government's former approach to carbon valuation. Nor is it necessary to use the SCC to set overall targets. Faced with uncertain but potentially catastrophic worst-case climate risks, the natural tendency is to set a precautionary target that implies a tolerably low level of risk, and then to seek the least-cost strategy for reaching that target.

This common-sense approach, which is one element of real-world policy deliberation, is also endorsed by recent economic analysis. In a mathematically sophisticated treatment of climate risk, Martin Weitzman has demonstrated that with potentially unlimited exposure to worst-case outcomes, it is impossible or irrelevant to estimate the expected

---

[1] Frank Ackerman, Elizabeth A. Stanton, Chris Hope, and Stephane Alberth, "Did the Stern Review underestimate US and global climate damages?", *Energy Policy*, forthcoming 2009, available on-line at doi:10.1016/j.enpol.2009.03.011

[2] See in particular Frank Ackerman and Lisa Heinzerling, *Priceless: On Knowing the Price of Everything and the Value of Nothing* (Washington DC: Island Press, 2004).

[3] See Richard Tol, "The Social Cost of Carbon: Trends, Outliers and Catastrophes", published on-line in *Economics E-Journal* (2008), http://www.economics-ejournal.org/economics/journalarticles/2008-25/view. On the same webpage, Tol provides a spreadsheet of the 211 estimates used in his review. My own calculations from that spreadsheet show that for the 99 observations with a reported discount rate and a positive SCC, the discount rate explains 46 percent of the variance in the logarithm of SCC. (Logarithms are used because SCC estimates span several orders of magnitude.) For the 169 observations with a reported pure rate of time preference and positive SCC, the pure rate of time preference explains 42 percent of the variance in the logarithm of SCC.

value of climate damages. Weitzman shows that the inescapable limits on scientific prediction, and the lack of a known upper bound on climate risks, imply that the true SCC is, technically speaking, infinite.[4]

Assuming that science-based policy deliberation has set a sensible target for emission reduction, therefore, it is appropriate to turn to MAC-based calculation of least-cost strategies for achieving the targeted level of reduction.

2. Interpreting abatement costs

MAC calculations are more solidly based than SCC estimates, in several respects. MAC projections are estimates of actual economic costs for marketed goods and services; they do not require monetary valuation of priceless human and natural values. Each component of a MAC curve reflects the costs for a single technology or policy innovation; modeling of global impacts is not required to get the costs right. The time spans involved are shorter than for damage calculations, lessening (although not eliminating) the role of the discount rate.

However, as attention appropriately focuses on the MAC curve, it is clear that there are important uncertainties surrounding the shape of the curve, and hence the carbon price implied to reach any given target. The draft illustrates this point by presenting alternate versions of the MAC curve, with the associated carbon prices required to reach the target for the non-traded sectors. While this is a useful concept, some of the uncertainties discussed in the draft, such as the actual potential for carbon reduction in transportation, and in agriculture, appear to be amenable to additional research. In view of the large effect of these uncertainties on the MAC-based cost of carbon, additional research to narrow the range of uncertainties should be a priority.

This section offers three principal comments on the uncertainties surrounding the MAC curve:

- oil prices will have an enormous effect on MAC estimates;
- discount rates and related assumptions remain important in MAC calculations, and should be explicitly discussed; and
- the draft's skepticism surrounding non-CO2 abatement is unwarranted; these are some of the lowest-cost opportunities for emission reduction.

Oil prices are mentioned only once in the draft, as the subject of an additional sensitivity analysis that has yet to be performed on MAC estimates. That sensitivity analysis could prove to be more important than the alternate versions of the MAC curve presented in the current draft. Moreover, the uncertainty surrounding the future price of oil is inescapable; it cannot be much reduced by additional research (unlike, for instance, the difference between estimates of the potential for transport emission reduction discussed in the draft).

---

[4] Martin Weitzman, "On Modeling and Interpreting the Economics of Catastrophic Climate Change," *Review of Economics and Statistics* (2009), **91**(1):1-19.

MAC estimates are incremental costs, representing the change in cost over current conditions, or "business-as-usual": for example, the cost of reducing emissions by buying a fuel-efficient automobile is not the full cost of the more efficient car, but only the increase in costs over a comparable car of average efficiency, net of the additional savings on fuel costs. This is why it is possible to have negative-cost abatement: the fuel savings can outweigh the additional investment or program cost.

This is also why the price of oil is so central to the MAC. (The same analysis applies equally to other fossil fuels, although coal and gas prices are generally less volatile than oil prices; to simplify the exposition, I restrict this discussion to oil prices.) At a low enough oil price, there would be no negative-cost opportunities to reduce emissions, since fuel savings would be too small to offset the costs of any new measures. At a high enough oil price, on the other hand, virtually anything that reduces fuel use and emissions could have negative net costs. Thus there is no such thing as a MAC estimate independent of an oil price forecast; the explicit or implicit assumption about the price of oil determines the level of the MAC curve. IPCC discussions of "no regrets" options, or negative-cost savings opportunities, have often identified the price of oil at which a certain level of no-regrets savings can be achieved. The MAC curves used for policy analysis should likewise be labeled with the assumed price of oil, and accompanied by sensitivity analysis showing the effects of changes in oil prices.

The astonishing volatility of oil prices, which have been both above \$120 and below \$40 per barrel within the last 12 months, poses a challenge for MAC-based calculations. For any emission-reducing measure, roughly speaking,

Marginal abatement cost = investment or program cost – avoided energy cost

If investment or program cost is a fixed amount, independent of energy prices, then the sum of (avoided energy costs + marginal abatement costs) is a constant. This is both an analytical issue, highlighting the importance of selecting the best available energy price forecast, and a policy question: should the cost of carbon for policy appraisal move inversely with actual changes in the price of energy? If so, how frequently should it be changed? Forecasts of long-term energy prices are central to MAC estimates; as a result, revisions of those price forecasts may entail revision of MAC curves and the MAC-based cost of carbon.

The switch from SCC to MAC makes discount rates less important, but it does not banish them from the analysis. MAC curves rest on calculations of the net cost of carbon-reducing measures, typically an up-front cost net of savings over some period of time. The well-known McKinsey curves calculate life-cycle costs for energy-saving innovations. In this context, the length of the life cycle, and the discount rate if any over that life cycle, become crucial. Vehicles and major appliances will be in use, affecting energy choices for more than a decade; building shell improvements and power plant investments will last for many decades. What is the appropriate time horizon and discount rate for computing the savings from the longer-lived measures in the MAC curve? Addressing this question requires an unpacking of the individual estimates in the MAC curve.

Ideally, it should be possible to present MAC curves based on different assumptions about life-cycle length and discount rates. A MAC curve seen from a private investor's perspective might use a market interest rate for discounting and a moderately short time horizon for calculating energy savings. A MAC curve seen from society's perspective might use a low discount rate and a long time horizon. The contrast between the two would be instructive for public policy: the investor's version shows the likely response that could be achieved through better information and other low-cost steps to remove minor market failures; the gap between the short-term private and long-term societal versions of the MAC curve reveals areas where more active policy initiatives may be needed. According to numerous studies, there are negative-cost opportunities for emission reduction, even from the private investor's perspective; however, there could also be a significant category of emission reductions that have positive cost to private investors but negative costs from a long-run societal perspective.

While the range of sensitivity analysis on the MAC curve could be expanded to address oil price uncertainty and discount rate effects, some other uncertainties could be reduced. The uncertainty about non-CO2 emission reduction potential, expressed in the draft, seems like a solvable empirical problem. Many studies have found that reduction of non-CO2 gases offers low-cost opportunities for abatement. Agriculture and waste management are less technologically advanced than some sectors of the economy; it is not surprising that modest changes in techniques and practices can reduce the emissions from these two sectors. Use of nitrogen fertilizers, better manure management, and changes in feeding practices can affect agricultural emissions; methane capture at landfills eliminates much of the waste emissions. These are attractive places to start, especially when compared to the difficulty of reducing emissions in many other areas. Rather than expressing uncertainty about the reliability of savings in these areas (with a noticeable effect on the estimated cost of carbon), it would be appropriate to put a priority on research and policy development to understand and expand the carbon reduction achieved in agriculture and waste management.

3. The price gap between traded and non-traded sectors

The most startling numerical result in the draft is the size of the initial gap between the carbon prices for the traded sectors (those subject to the EU ETS) and the rest of the economy, £29 vs. £70 per tonne of CO2-equivalent. The explanation of the two prices is entirely logical, but suggests that the ETS cap is much higher than the targets used to set carbon prices for the non-traded sectors. The projected evolution of prices over the next 20 years shows the non-traded price remaining roughly constant, waiting for the traded price to catch up, before they begin rising together after 2030. While it is hard to fault the exposition of either the traded or the non-traded price, something must be wrong with this picture.

The gap between the two prices means that UK policy is not able to use some of the least-cost measures for emission reduction – namely, opportunities for reduction at more than £29 but less than £70 in the traded sectors. A likely political response is sheer disbelief, and refusal to accept a gap of that magnitude; if the ETS cap and resulting price are unchangeable, the result could be a politically driven reduction in the non-traded price to something much closer to the traded sectors' level.

It is possible that the legal situation prevents any remedies for this disparity; I do not pretend to any expertise in British or European law. But if there is any way to mobilize additional savings in the traded sector, beyond £29 per tonne, it may reduce both costs and political opposition to the carbon policy. Fears of "leakage" – of production moving to other jurisdictions with lower carbon prices – are often exaggerated; businesses rarely move in response to small changes in environmental regulatory costs, and electric power generation cannot move out of the UK at all.

In the U.S., informal discussion has begun among state policymakers who anticipate that states may want to impose a stricter cap on carbon emissions than allowed by a projected national cap. Some of the options under consideration (the relationship to American law is still unclear) include a tax or surcharge on national carbon allowances; creation of parallel state allowances, which would be required in addition to the national ones; and perhaps most promising, a requirement that firms operating in a particular state buy or own additional allowances, such as 1.25 allowances per tonne of emissions. In the latter proposal, the additional 0.25 allowance per tonne would be retired unused as a cost of doing business in the state, effectively raising the state price of carbon in the traded sectors to 1.25 times the national price. This would create an incentive to adopt abatement measures with costs up to 1.25 times the national allowance price.

4. Response to questions asked of peer reviewers

Many of these answers are implied in the above discussion; they are listed again here for convenience.

Q1: Do you agree with the arguments for using estimates of marginal abatement costs to value carbon in individual policy appraisal, once targets have been set?

Yes, I am strongly in favor of this approach, as explained in section 1 above.

Q2: Do you agree that social damage cost estimates should still be used along with other evidence to when setting emissions reductions targets?

No, I believe that social damage cost estimates are not adequate to this task. See section 1 above.

Q3: Do you agree that two carbon values are needed for policy appraisal in the short-run (up to 2020), one to assess EU Allowance impacts in the traded sector and another to assess carbon impacts in the non-traded sector?

Yes, although the gap between the two is disturbingly large, a practical political problem that may need to be addressed to make the policy credible; see section 3 above.

Q4: Do you agree with the approach taken to setting carbon values in the long term (post 2030), based on modelled estimates of global abatement costs and emissions reductions trajectories consistent with global GHG stabilisation goals?

Yes, this seems like the appropriate way to set long-term prices, consistent with the MAC-based approach to pricing.

Q5: How does the proposed approach compare to any other international approaches you are aware of?

I am not aware of any similar international approaches; see my comments on U.S. state policy discussion in section 3.

Q6: Do you agree that a modelling approach rather than observed market prices should be used as a proxy for the short term estimates of the traded price of carbon?

This seems reasonable; it does not address my larger concern about the gap between traded and non-traded prices in the short run.

Q7: A particular set of issues arises with the calculation of the short term non traded price of carbon, given the large disparity in 2020 with estimates of the traded price of carbon and the variation in the price under different scenarios:

> *i) Do you agree with the approach taken to dealing with uncertainty in the modelling of abatement costs in the non-traded sector in the short run, i.e. looking at a range of feasible scenarios to produce a central estimate and range?*

Yes, although I think that a different set of uncertainties should be highlighted, as discussed in section 2 above.

> *ii) Does the evidence suggesting a large disparity between traded and non-traded values seem robust?*

In a sense; it is robust only because there are apparently inconsistent assumptions about carbon reduction goals in the ETS cap versus the targets for the UK that drive the non-traded sector price. With consistent carbon reduction goals, the disparity should largely or entirely disappear. The inconsistency of goals, not the modeling techniques, is the important problem.

> *iii) Do you agree with the approach taken to deal with the potential for endogeneity in the setting of the non traded price, through basing the analysis on estimates of feasible technical potential (as opposed to policies) in the abatement curves and through the use of a variety*

of scenarios? And are you aware of any precedents for tackling similar problems elsewhere in cost-benefit analysis?

I am not sure why this is thought to be unique. It seems consistent with many cost-benefit analyses of new policies and technologies. I suspect that analogues could be found in US EPA analyses of the costs and benefits of the Clean Air Act, for instance.

iv) Do you agree that air quality impacts and other ancillary impacts should be factored into the abatement costs?

Yes, if robust estimates for other air quality impacts are available, they should be included. Reduction of fossil fuel combustion (and agricultural and waste management emissions as well) often entails reduction in other, better-studied and regulated emissions; when monetary values for these emissions are known, it is perfectly reasonable to include these "co-benefits" in calculating marginal abatement costs for greenhouse gases.

Do you think that adjustment to abatement costs should be made to reflect policy costs and if so, how should that adjustment be made?

In theory, policy costs should be included; however, as suggested by the McKinsey report quoted in the draft, policy costs are probably small, and very difficult to estimate. It is important to *avoid* the inference made by some economists that policy costs must exist, and must be large, because negative-cost savings are known to be impossible a priori. Needless to say, this does not represent a measurement of policy costs; rather, it is a rejection of empirical information about abatement costs, driven by abstract doctrine.

If the McKinsey estimate of policy costs is accurate, then it is only a second-order change in the carbon price to include or exclude them; it is a much smaller effect than the uncertainties discussed in the draft and in my comments.