National Research Council: Committee on Assessing Approaches to Updating the Social Cost of Carbon

Notes from Presentation to Session 2, Part 2 Current State of Evidence and Approaches, Future Research Needs

John Reilly

It is not clear to me that the Social Cost of Carbon is a particularly useful construct with which to guide policy. The theoretical foundation for the construct rests on weak grounds. Even if we accept (or ignore) the theoretical flaws and proceed the practical and "philosophical" issues that must be resolved are many—and again its not clear to me that theory or logic is a positive guide to these decisions. Then, even if we finally resolve these issues through some process. (A vote among economists, a ruling by the Office of Management and Budget, a thrashing out by this committee of very specific guidance on how to proceed?), the empirical challenges to meaningfully estimate (improve on the current estimate of) a social cost of carbon are immense. I do not really see a way forward that will meaningfully improve estimates we have now (in the one to three year period given as guidance to this workshop in the agenda). I'm sure all here are fully aware of the above issues, but just to quickly review some of these so that we are all on the same page.

The Social Cost of Carbon concept relies on the existence of a social welfare function. Three articles by Paul Samuelson spanning the mid-1950's to early 1960's demonstrated graphically and mathematically that a well-defined social welfare function does not exist (e.g. Samuelson, 1954). I reviewed and discussed this literature in my Ph. D. dissertation (Reilly, 1983). The Samuelson argument: The social welfare function aggregated across individuals reflects their current income. A policy that changes income (the very policy for which the welfare construction is done) will change the welfare function. So the policy welfare function might lead to a different social cost of carbon than the one on which the policy was determined. Of course economists have ignored those arguments and have continued to use social welfare functions. Perhaps there is something useful in the exercise even if the theoretical foundation is non-existent. Maybe it gives some guidance on how to organize an investigation, or bring together information that is in some way useful.

Of course the entire discounting issue is crucial to estimating damages, and has been the subject of many workshops over the years. At one of these the question was raised: Was it really necessary to discount to get a single net present value, or might decision-makers (and the public) be better informed by simply showing the full time path of damages (I suppose under different circumstances of mitigation continuing or not). This could then feed into a democratically debated solution about what to do now, the result of which would be an implicit marginal abatement cost, or maybe cap trade system with observable price, or a carbon tax that reflected a "willingness to pay". One might compare it to a social cost of carbon constructed in some other way—with models as has been done so far—and perhaps that would lead to some head scratching about why they were different (or coincidentally the same). I think there is, however, an inevitable conflict between democratic solutions (one person one vote) and economic valuation where individual weights are their "willingness (and ability) to pay." Of course perhaps that conflict may be less than we imagine given the role of money in our democratic process.

So the basic theoretical foundation of the social cost of carbon rests on there being a stable global social welfare function, which there is not. But the charge here is to move on.

Accepting the existence of a global social welfare function, the practical and "philosophical" challenges are that there is no global government institution, or as I at least in my training the term "benevolent dictator" was used who was going to evaluate and enforce the premise of an optimal cost benefit solution. We of course introduce the concept of a Pareto improving change, where, in principle everyone could be better off if the right compensation scheme could be devised. This allows us to avoid the need for interpersonal comparison of welfare. Within a nation with a strong central government, one can perhaps believe that over time with many decisions it all evens out. This seems a harder case to make given our global governance structure. The practical: Should the US estimate a social cost of carbon based only damages to the US and its interests (and if so where do those begin and end)? Should it take the perspective of a benevolent dictator looking out for the world? Whose valuation of impacts should we use—or whose judgment about who's values should we use? The potential disappearance of small island nations is an example. There are a number of these issues. The Stern Review discussed many of these at length (Stern, 2007)—it came up with a number none-the-less. It was strongly criticized especially for the discounting (e.g. Nordhaus, 2007) but the theoretical discussion of the difficulties around the concept of a global damage function (which implies some global welfare function by which damages are valued) was carefully constructed, reviewing conventional neoclassical economics.

But lets assume those issues away and move on to the really practical aspects of coming up with a social cost of carbon. As economists this is an issue of valuation. Tell me what happens, and with my economic tool box I will put a value on that impact. Assuming away the theoretical issues I raised valuation, in its simple application is not a hard problem. Some years ago, a colleague of mine was involving on preparing testimony on the value of salmon in case in the Pacific Northwest about dams and what they were doing to the natural salmon population, on which native Americans in the region depended on as part of long cultural association with the region. I don't even recall which side he was preparing testimony for but the judge, on hearing a valuation argument, provided an answer of

$7.99 a pound based on his visit to the supermarket the night before. If it is a pure loss of a market good—drop in the yield of corn, damage to a roadway that can be repaired—economic methods are very good at telling us the lost value of a marginal change in the amount of the good available. Unfortunately, climate change of 3, 4, 6, 8 degrees C is not "marginal." We won't just lose one beach, we will lose all of them (maybe some new ones will form). We aren't talking about the value of one ski resort (assuming all others remain unchanged). They may all disappear. Economic valuation methods are good for marginal changes for market goods with established prices—but climate change is non-marginal and affects lots of things which may be part of our culture that we (or someone) values in ways we can't easily estimate or describe.

Then we still do not really know what the physical impacts are. We have good projections of climate change at the level needed for looking at impacts that happen at specific sites. We don't have the geographical resolution, nor the resolution of extreme events. The climate modeling community, steeped in science has proceeded in good scientific method. This is to accept the null hypothesis unless demonstrate with a high degree of confidence that an alternative hypothesis should be used. This means often that there deep suspicions of what might go wrong and how the climate might change, does not find its way into a climate scenario until that suspicion can be tested, proven, and the mechanism by which it operates fully described. The IPCC AR4, WG1 report which quantitatively estimated sea level in the future but excluding large ice sheets is an example. From a science standpoint, without good understand of the process they could not provide an estimate. From a decision science standpoint, their estimate of sea level increase was almost certainly biased low. There are many of those types of issues lurking throughout climate science. And there are many things we simply don't understand about how the climate system works—the "black swans" or "unknown unknowns." This is the Weitzman argument for fat tails (Weitzman, 2009).

A decade or more ago, I participated in a similar NAS meeting as invited outside presenter on abrupt climate change. The chance of abrupt climate change is one of those mostly unknown, unknowns. We think it possible, there seems to be records of abrupt climate change that can be ascertained from ice cores or sediments, or something. But we don't know what caused them, very much about how big they were, or how fast. We don't know whether global warming is making them more or less likely, and we really don't have sense of the timing. Moreover, I raised the issue in that paper/presentation that abrupt climate change that is relevant for human activity, need not register as an abrupt change in the global mean surface temperature. If the climates that support much of our food production in breadbasket areas of the world, simply shifted west or east a few hundred miles, this might well turn current breadbaskets into deserts or flood zones. Could this happen quickly—I don't think climate models are in state where they can tell us. And, such shifts would not occur randomly through time. One region here this year, a different region 30 years later… It is a climate system tied together in a general circulation that goes around the globe. So if something big happens in North America, there probably is something big in Europe and Asia at the same time.

Then apart from fat tails because of deep uncertainty, even if we try to characterize the distribution of future outcomes based on our known understanding of the earth system as represented in existing earth system models, we know that distribution is wide. It is impossible to simulate hundreds or thousands of simulations of complex climate models that have the resolution vaguely close to what we need. Our own work in the MIT Joint Program on the Science and Policy of Global Change has focused on trying to do that at some level (e.g. Sokolov, et al., 2009; Webster et al., 2012; Reilly, et al., 2013; Fant et al., 2014). We know adaptation is important, but we don't know what we are adapting to. We require a stochastic dynamic representation of the adaptation decision problem, wrapped around a highly resolve Integrated Assessment Model or Integrated Earth System model, in which everything is changing simultaneously.

Around the time I participated in the aforementioned NAS study on abrupt climate change, I had just read the book "When Genius Failed." It was a description of the disastrous end of the firm, Long-Term Capital Management, that had been built on idea of derivatives for which advisors to firm had won a Nobel prize in economics. The diagnosis: (1) The full uncertainty in possible outcomes was not taken into account—the distributions were based on just a few years of data where nothing bad happened. (2) The estimation assumed normal distributions and so there were no fat tails. (3) Risks in different markets were assumed to be independent. (4) The firm assumed it was a small actor in the market and so good get out at any time (but had been so successful that many financial entities were watching its every move and so if they tried to exit, everyone would rush through the door.)

That is a perfect description of climate change and our methods for assessing damages. (1) We have not characterized well the full range of possible outcomes. (2) Unknown, unknowns make fat tails likely but they mostly don't show up in our modeling, (3) It is a global system, with global climate change affecting everything simultaneously, and so the very thing we imagine we would use to adapt to an impact, may itself be gone. (4) Our analysis is based on marginal valuation, when it's a non-marginal change.

What could really change our understanding of the social cost of carbon?

(1) Given discounting moderate impacts, several decades to centuries won't have much effect. If we thought we could better predict the next 5 to 30 years, and understand what was happening or might happen, we might find much worse impacts than we currently find with out methods that wash out extreme events, smear out geographical richness of impacts through coarse representation of climate, probably miss important channels of impact, and assume idealize adaptation. Most of our climate impact work looks out 50 or 100 years and extrapolate back to present in some simple way. We are now seeing large number

of deaths in extreme heat, witnessing extreme flooding and property destruction, and seeing other impacts scientists are better able to link to climate change.

(2) While there will always be “unknown unknowns” more effort to work with scientists to understand suspected bad outcomes either because of high climate response, or just because of patterns of climate change (such as e.g. abrupt change). If these events are large scale and catastrophic, then valuing them may be largely irrelevant—“catastrophic” may tell us all we need to know about the value—its big. But many of these extreme outcomes may be not globally catastrophic but if we add up enough it may come up to a big number—even if discounted from the relatively distant future. So serious estimation of the risk of these occurring and the economic loss could help.

(3) Simple econometric studies focused on one part of one sector in one region and based on marginal valuation are the start of something but hardly the final answer. We need to look at these impacts in a system that can consider non-marginal change, and sectoral and regional interactions.

(4) The real problem of adaptation in a very noisy climate system needs to be seriously investigated. This likely means added cost.

(5) While the social cost of carbon is at its heart a valuation exercise, which calls on the economists tool box, it is very dangerous to take climate model outcomes and scenarios as “representative” of future climate in a way that is actually useful for a decision science question, which the social cost of carbon estimate is. So, economics research on this needs to join closely with the science community.

Fant, C., C.A. Schlosser, X. Gao, K. Strzepek and J. Reilly, 2014: A Framework for Analysis of the Uncertainty of Socioeconomic Growth and Climate Change on the Risk of Water Stress: a Case Study in Asia, MIT Joint Program on the Science and Policy of Global Change, Report No. 269,

Nordhaus, W.D., 2007. Review: A Review of the "Stern Review on the Economics of Climate Change” *Journal of Economic Literature* **45(3**): 686-702.

Reilly, J. 1983: *An Economic Analysis of Global Energy and Carbon Dioxide Emissions*, University of Pennsylvania.

Reilly, J.M., S. Paltsev, K. Strzepek, N.E. Selin, Y. Cai, K.-M. Nam, E. Monier, S. Dutkiewicz, J. Scott, M. Webster and A. Sokolov, 2013. Valuing climate impacts in integrated assessment models: the MIT IGSM, *Climatic Change* online first, doi:10.1007/s10584-012-0635-x.

Samuelson, P. 1954; The theory of public expenditure, *Review of Economics and Statistics* **36**: 387-389.

Sokolov, A.P., P.H. Stone, C.E. Forest, R.G. Prinn, M.C. Sarofim, M. Webster, S. Paltsev, C.A. Schlosser, D. Kicklighter, S. Dutkiewicz, J. Reilly, C. Wang, B. Felzer, J. Melillo, H.D. Jacoby, 2009. Probabilistic Forecast for 21st Century Climate Based on Uncertainties in Emissions (without Policy) and Climate Parameters, *Journal of Climate*, **22(19)**: 5175-5204

Stern, N., 2007. *The Economics of Climate Change: The Stern Review*. Cambridge and NewYork: Cambridge University Press.

Webster, M., A.P. Sokolov, J.M. Reilly, C.E. Forest, S. Paltsev, A. Schlosser, C. Wang, D. Kicklighter, M. Sarofim, J. Melillo, R.G. Prinn and H.D. Jacoby, 2012. Analysis of climate policy targets under uncertainty, *Climatic Change*, **112(3-4)** 569–583.

Weitzman, M., 2009: On modeling and interpreting the economics of catastrophic climate change. *Review of Economics and Statistics* **91**: 1–19