False confidences: forecasting errors and emission caps in CO_2 trading systems

Michael Grubb*, Federico Ferrario

Faculty of Economics, Cambridge University, Sidgwick Avenue, Cambridge CB3 9DD, UK

Abstract

This Commentary sets out four lines of evidence to argue both that emission forecasts are intrinsically *uncertain*, and that there is clear evidence of *projection inflation* in the forecasts of sector emissions used to underpin the setting of sector caps in emission trading systems. From a limited evidence base, we conclude that uncertainty is at least ±2%/year, overlaying an upward bias (projection inflation) on the order of 1%/year, cumulative. The Commentary concludes that this has important implications both for allocation approaches, and for some other design elements in the EU ETS. Forecasting uncertainty is not an inconvenience which is best ignored, but a fundamental fact that must be accommodated in the future design of the EU ETS and other CO_2 emission trading schemes.

Keywords: Emissions trading; Forecasting; Emission projections; EU ETS

Introduction

> Energy forecasting was invented to make economic forecasting look good.
>
> Anon.

So wrote one cynic, looking at the history of energy forecasts from the 1970s in the aftermath of the oil shocks. Yet understanding the 'science of forecasting' has assumed a wholly new significance in the context of allocating allowances for emissions trading. If allocation risks being the 'Achilles' Heel' of emissions trading, then projections of sector emissions form the protective hand responsible for a false sense of confidence.[1]

The EU ETS has hinged upon imposing cutbacks of just a few per cent, usually relative to projections of 'business as usual' emissions. The politicians who make the ultimate decisions around allocation tend to trust the numbers provided by their technical experts, who put long hours into economic modelling of projections and, increasingly, discussion with representatives of the sectors affected by allocation decisions. When the cutbacks are small relative to projections,

* Corresponding author. Tel.: +44-1223-335290; fax: +44-1223-335299
E-mail address: michael.grubb@econ.cam.ac.uk

the entire carbon market then hinges upon the accuracy of these projections. It is, consequently, remarkable how little critical analysis there has been of the reliability of such forecasts, despite questions being increasingly raised.

This Commentary sets out four lines of evidence to shed light on forecasting errors, and concludes by discussing implications for the future development of CO_2 cap-and-trade systems.

Scenario projections

Scenario projections are used because of acknowledged uncertainty about the future. They shed light on the range of uncertainty that forecasters themselves acknowledge to be plausible, and imply a wide range of possible emission levels. For the long-term global scenarios, the IPCC Third Assessment (IPCC, 2001) charted 'non-mitigation' scenarios ranging from under 10 to over 30 GtC/year emissions by mid-century – corresponding to average growth rates in the range 0.5–3%/year. The Fourth Assessment (IPCC, 2007) scarcely narrows the range. Moreover, by definition, the more disaggregated the focus – by sector, region and/or time-span – the wider the range of uncertainty within the global envelope.

At a national level, on time periods relevant to the EU ETS phase-II allocations, a recent study by the UK Department of Trade and Industry (DTI, 2006a) warned that UK national emission forecasts may vary by ±6% just because of modelling and statistical uncertainties. In this Special Issue, Neuhoff et al. (2006) show how, by adopting the four different fossil fuel price scenarios in the DTI study and accounting for a range (±0.75%) of expected economic growth, projections for EU ETS sectors for phase II (2008–2012) may vary as much as ±9% of the central estimate; the central 60% of scenarios still spanned an uncertainty range of 6%, over just 5 years. Such ranges would, of course, considerably increase if extended to the outlier scenarios and other modelling studies.

Other factors add to uncertainties, including questions about the reliability of historic data and the expected impact of other policies and measures to reduce emissions. The verified emissions data for the EU ETS in 2005 represent the best data available, and imply that the less accurate historic emissions used so far would have created additional (and biased) errors in the projections based on them.[2] Various policies and measures adopted by EU Member States in their efforts to comply with the Kyoto Protocol include sectors covered by the EU ETS, with the degree of response to these measures introducing additional uncertainties.

Even based on scenarios alone, this suggests a typical forecasting uncertainty of at least ±2%/year cumulative. The observation that the future is not certain due to uncertainty in underlying parameters is, of course, no surprise – though it is still not obvious how that is factored into allocation processes. More troubling is the possibility of *bias* in projections, on which other approaches to the issue shed light.

Statistical analyses of past forecasts

The most rigorous approach to assessing forecasting errors would be to analyse statistically the accuracy of past forecasts. There is a well-known history stemming from the 1970s debates about energy – both projections of oil depletion in the face of galloping demand, and forecasts of electricity consumption used to justify the planning of new power stations. In both cases, the 'energy establishment' was widely criticized for forecasts that turned out to be far too high. Of course, with hindsight there were plausible reasons: the oil world underestimated both the demand and supply responses to the price shocks; and electricity analyses 'projected the past' in terms of economic growth rates that with hindsight were clearly unrealistic, as well as underestimating the scope for improved efficiency.

The general aura of 'projection inflation' from that era is only partly true. As energy prices, markets and macroeconomic conditions stabilized, and energy economists absorbed the lessons, energy demand forecasts became more modest.

In one of the most interesting contributions of the time, Baumgartner and Midttun (1987) collected analyses of eight countries, together with the global forecasts by IIASA, to establish overwhelming evidence that 'largely because of the importance of forecasters in the policy process, they are subject to a variety of influences which combine to prevent their forecasts from being objective'. For example, in Denmark, the dominant coalition had in the 1970s reflected industrial interests which wanted, and projected, big growth. Not long afterwards, when a far more environmentally inclined government came to power, energy projections dropped dramatically in favour of 'low-growth' projections. The 'politics of energy forecasting' is a theme to which we return briefly below.

In a bid to make energy projections more objective, in the USA responsibility rests with the Energy Information Administration, a specialist arm of the US Department of Energy, which has over many years built up a track record of readily available, annually published projections. Unlike the sporadic, *ad hoc*, and sometimes unpublished projections in most other countries, the EIA series facilitate rigorous statistical analysis. A recent study (Winebrake and Sakva, 2006) reports:

> Low errors for total energy consumption are concealing much larger sectoral errors that cancel each other out when aggregated. For example, 5-year forecasts made between 1982 and 1998 demonstrate a mean percentage error for total energy consumption of 0.1%. Yet, this hides the fact that the industrial sector was overestimated by an average of 5.9%, and the transportation sector was underestimated by an average of 4.5%. We also find no evidence that forecasts within each sector have improved over the two decades studied here.

Unfortunately for schemes like the EU ETS, it is precisely the industrial sector, not the national total, that matters. And a consistent upward bias error of more than 5% in as many years, if translated into the European context, would imply huge problems where the regulatory system is seeking emission cutbacks from projections that are not much bigger than the systemic bias identified in US industrial energy forecasts.

The remarkable fact is that, despite a search in both the literature and in web-based data, we found neither published analyses nor a readily analysable data-set that would enable a similar analysis to be carried out for European countries.

Understanding the process

The third line of evidence is to understand the process of constructing official emission forecasts. The starting point is economic or sector modelling, which needs to be informed by projections of underlying driving forces (such as economic or sector output growth) and other influential parameters (such as fuel prices), together with estimates of response functions (notably, elasticities of fuel consumption with respect to changes in the various input parameters). Most of these can now be estimated econometrically, hence 'using the past' to project the future.

Such projections err to the extent either that input assumptions on driving forces prove wrong, or if future responses are not a continuation of past patterns. Energy price forecasting is notoriously uncertain, almost the only systematic feature being a tendency to place too much emphasis upon recent trends, and related 'groupthink' problems as the world cycles between high and low energy price periods. The request to revise UK allocations upwards for phase I of the EU ETS reflected to a large degree the impact of sharply rising gas prices, which drove generators back towards coal, sharply increasing emissions.

Economic growth and sector output projections are potentially more liable to systematic error, as there are strong pressures towards optimism. No government likes to predict a gloomy economic future, or imply that it will mismanage the economy. No company raises capital, justifies a new project, or energizes its workforce, by proclaiming a future of decline. Indeed, capitalism thrives on the optimism of those in the market. Yet not all can be winners: competitive markets are all about weeding out those whose optimism proves misplaced.

This introduces a paradox into forecasting. The more disaggregated the level of forecasting – the greater the level of sectoral and even plant-level detail – the more information is at hand, and the more precisely sector forecasts can match the detailed plans and projections of the actual players in the market. Such rich information is lost in aggregated, economic model-driven, forecasts. Yet only aggregated forecasts can capture the statistical fact that not all the individual plans and hopes will come to pass.

This paradox gives rise to real procedural dilemmas. For example, to support its allocation process, the UK put its model-based projections out for consultation, and established an expert advisory group to strengthen the independence of its response to industry representations. Unquestionably consultation reduces the risk of significant underprojection, as industry will have a strong incentive to avoid that. Yet the converse is lacking: if projections are used to inform allocation, what industry would interject to argue that its projected emissions are too high? The level of detailed knowledge required – and the intrinsic uncertainty – makes it impossible for any degree of independent oversight to be certain about the accuracy and objectivity of all evidence for revising sector projections.

In the case of the UK, this effect was balanced by maintaining an aggregated model-based cap across ETS sectors, derived from Treasury projections of economic growth, with the detailed consultations being used to inform the sector distribution; the overall process did result in some modest downward revisions (by 4% for the non-large electricity producers in aggregate; DTI, 2006b). Few other Member States have adopted such an extensive process, and considerable uncertainties remain.

The history of allocation negotiations

This brings us to the final line of evidence, namely, the actual *ex-post* history from cases where governments have sought to negotiate allocations of CO_2 emissions, in one form or another. The UK offers two clear examples.

The first concerns the pilot UK emissions trading scheme, a 'bid-in' scheme in which 32 companies were paid a total of £215m to accept binding caps (established though an auction process). Collectively, participants agreed targets to deliver emission reductions on a linear increase to 3.8 $MtCO_2e$ by 2006, a reduction of around 14% relative to collective baseline of 27.8 $MtCO_2e$. In practice, trading prices peaked after a few months, and then fell towards zero as it became increasingly clear that the market was in surplus. The official assessment of the scheme (National

Audit Office, 2004) studied four of the biggest participants (which accounted for more than 50% of the incentive funding) and reported:

> In 2002 their emissions were 3.78 million tonnes below their baselines, nine times the target of 0.42 million tonnes ... approximately 66% of the reductions reported by these four companies is attributed to the Scheme while an estimated 34 per cent (1.28 million tonnes) is not.

By the end of the second year, 6.5 $MtCO_2e$ of surplus allowances were already banked forward (NERA Economic Consulting, 2004). Not surprisingly, the price had dwindled close to zero, and any discussion of extending the scheme was shelved as the EU ETS emerged.

On a larger scale, the UK Climate Change Agreements gave 44 industrial sectors discounts from the UK Climate Change Levy if they met a set of negotiated targets for emission reduction, defined biannually through to 2010. Despite the enormous effort expended in negotiating the original targets, the second target period assessment report is a testimony to the inherent difficulties:

> ... [since 2000] there have been widespread structural changes in UK industry, changes to products because of market forces and entrants and exits in many sectors ... the assumptions of growth and energy prices on which the original BAU forecasts were made are now outdated and of limited relevance.

This, it should be remembered, refers to a period generally hailed as being in the middle of the UK's longest period of sustained, stable, low-inflation economic growth and, in its early years, relatively stable fuel prices. Excluding the steel sector, in target period 1 (2002) the savings were about 40% greater than the targeted savings relative to base year; for the second period, the savings were more than double the target. Exceedence in the steel sector was much bigger still. As a result, the review process resulted in a tightening of targets for the final three target periods (compared to the original agreements) in all but four of the sectors.

Finally, the biggest test of all was the 2005 verification data on the EU ETS. The most recent and detailed analyses by Kettner et al. (2007) confirm that allowances issued for 2006 exceeded verified emissions by about 100 $MtCO_2$, or about 5% of the total.3 Moreover, sectoral analyses demonstrated a highly skewed distribution, in that the 'net short' element was restricted to about a 60 $MtCO_2$ shortfall in the power sector of EU15 countries (dominated by the UK), set against the much larger surplus from other sectors, and in all sectors in the Accession countries. The percentage surplus in non-power sectors ranged from 5% to 30%, or higher in some cases.

The analysis by Kettner et al. (2007) also presents evidence from the 2005 verification data that the bigger facilities were more able to lobby effectively for allocations that turned out to be inflated – within just a year or two.

Ellerman and Buchner (2006) discuss the difficulty of estimating how much of the surplus was due to abatement, concluding that there was unquestionably *some*, but that any estimate of how much is 'arbitrary and must remain so until better data and more careful assessments can be made'. They develop a 'net ratio' indicator to assess likelihood of over-allocation. Whilst the previous examples in this Commentary focused on the UK, the UK and Ireland show the *least* evidence of aggregate over-allocation (with a significant net shortfall); whereas at least half the Member States have ratios (+0.4 or higher) that could be interpreted as suggesting a general pattern of over-allocation across all sectors.

Note that all three examples share the fundamental impossibility of knowing the 'counterfactual' – what would have happened if the schemes had not been introduced. Thus, they all embody the impact of the schemes on incentives to reduce emissions, as well as the possibility of projection inflation. All attempts to assess retrospectively who or what was responsible conclude that it has been some hard-to-apportion balance of actual abatement with inflation of the original projections. The corollary of estimates that emphasize the importance of abatement compared to projection inflation, however, is that they imply that abatement must have been much easier than originally anticipated during the course of allocation negotiations.

Conclusions and implications

These four lines of evidence – scenario projection ranges, statistical analyses of past forecasts, the process of forecasting, and the evidence from previous allocation and target-setting efforts – point to two fundamental conclusions:

- There is an intrinsic uncertainty in forecasting energy consumption and emissions, probably exceeding ±2%/year cumulative, that no amount of analytic sophistication, whether model-based or procedural consultation, is likely to remove.
- Underlying this, there is evidence of systematic upward bias in industrial energy and emission forecasts, particularly (but not only) when these form the basis of setting sector emission targets or caps, probably on the order of +1%/year.

Obviously, certain kinds of fluctuations can result in bigger uncertainties and biases, particularly in the short term. The numbers here are probably reasonable minimum estimates for forecasts relevant to the timescales of allocation decisions, and yield for example an 'expected error' in the range of at least –5% to +15% in forecasts for industry sector emissions for the EU ETS phase-II allocations (based on approx. 5-year forecasts). This may sound large; yet it is modest compared to the range of sector-level gaps between allocations and emissions in the EU ETS phase-II 2005 data, just 2–3 years after the underlying projections were made.

At least four implications flow from this.

First, any emissions trading scheme in which sector cutbacks are estimated relative to emissions projections, and are less than or comparable to the intrinsic forecasting error, risks a considerable degree of shortfall or surplus at the sector level, even prior to abatement. In the EU ETS phase II, most countries have proposed this approach for sectors other than electricity – given underlying bias errors, the a *priori* likelihood is that most such sectors stand to gain irrespective of their pricing and abatement strategies.

Second, even *aggregate* phase-II cutbacks across all sectors and countries on the order of 5–10% relative to projections risk creating an intrinsically unstable market, since there will be no way of knowing for sure whether in reality the cutbacks during the period are actually extremely modest (if the 'business-as-usual' trend is below that forecast), or severe (if it is higher than forecast).

Third, future development of trading schemes must acknowledge the importance of both irreducible uncertainty and projection inflation, and be designed accordingly. Avoiding projection-based allocation – e.g. through benchmarking – would reduce the direct incentive to projection inflation (although indirect incentives could remain). But even with this, a stable market can only be ensured if the cutbacks are much bigger than the intrinsic projection uncertainties. If this is deemed economically too severe, other approaches must be considered – for example with consideration of measures explicitly designed to contain price uncertainty within bounds, such as mechanisms for price floors and ceilings (see, e.g., Hepburn et al., 2006). Even if free allocations are cut back in line with declared European goals for deep emission reductions – which is clearly not the case for most proposed phase-II allocations at the time of writing – such mechanisms may still have a valuable role to play in increasing confidence, both that the costs are containable, and that prices will support low-carbon investment irrespective of other considerations (such as CDM inflow).

Finally, projections can only conceivably improve if data are reliable, consistent and accessible to independent analysis. This is not the case at present. The relevant data in Europe are scattered, varied and frequently opaque, even when technically accessible. Improving data quality and comparability, and encouraging independent scrutiny over extended time periods, is essential if the EU ETS is to evolve into a more stable incentive for low-carbon investment.

Acknowledgements

We are grateful to numerous colleagues for the information and guidance that informed this Commentary.

Notes

1 The expression (used in relation to EU ETS phase-I allocations in Grubb et al., 2005) derives from the Greek myth in which Achilles' mother gave the baby godly protection by dipping him into the magical waters of the river Styx – but held on to his heels, which thus became his only vulnerable point.
2 As the 2005 Verified Emissions have been consistently below most expectations.
3 See Note 3 in Betz and Sato 2006, this issue.

References

Baumgartner, T., Midttun, A. (Eds), 1987. The Politics of Energy Forecasting: A Comparative Study of Energy Forecasting in Western Europe and North America. Oxford University Press, New York. Specific citation from review by R.J. Jones, International Journal of Forecasting Policy 5 (1989), 133–145.
Betz, R., Sato, M., 2006. Emissions trading: lessons learnt from the 1st phase of the EU ETS and prospects for the 2nd phase. Climate Policy 6(4), 351–359.
DTI, 2006a. UK Energy and CO_2 Emissions Projections: Updated Projections to 2020. Department of Trade and Industry, February 2006.
DTI, 2006b. EU ETS Phase II CO_2 Emission Projections: Revisions to Projections following consultations. Department of Trade and Industry, August 2006.
Ellerman, D., Buchner, B., 2006. A preliminary analysis of the EU ETS based on the 2005 emissions data. Nota di Lavoro 139.2006 [also downloadable through http://www.climate-strategies.org/].
Grubb, M., Azar, C., Persson, M., 2005. Allowance allocation in the European emissions trading system: a commentary. Climate Policy 5(1), 127–136.
Hepburn, C., Grubb, M., Neuhoff, K., Matthes, F., Tse, M., 2006. Auctioning of EU ETS phase II allowances: how and why? Climate Policy 6(1), 137–160.
IPCC, 2001. Climate Change 2001: Mitigation. Cambridge University Press, Cambridge, UK.
IPCC, 2007. Climate Change 2007: Mitigation. Cambridge University Press, Cambridge, UK (forthcoming).
Kettner, C., Kopp, A., Schleicher, S., Thenius, G., 2007. Stringency and distribution in the EU Emissions Trading Scheme – the 2005 Evidence. Climate Policy, forthcoming.
National Audit Office, 2004. The UK Emissions Trading Scheme: A New Way to Combat Climate Change. HMSO Report HC-517, NAO, London.
NERA Economic Consulting, 2004. Review of First and Second Years of the UK Emissions Trading Scheme. NERA, London.
Neuhoff, K., Ferrario, F., Grubb, M., Gabel, E., Keats, K., 2006. Emission projections 2008–2012 versus national allocation plans II. Climate Policy 6(4), 395–410.
Winebrake, J.J., Savka, D., 2006. An evaluation of errors in US energy forecasts: 1982–2003. Energy Policy 34, 3475–3483.