Peer Review of DECC's 2012 update to short-term traded carbon values

Dr William Blyth, Oxford Energy Associates, July 2012

Results of the POLES model used to inform DECC carbon price modelling suggests that the EU-ETS is now significantly over-supplied with allowances compared to business-as-usual demand, a fact which has changed since the DECC 2011 price estimates were produced. The balance of supply and demand is a key price driver in the market, but the scope of this review does not include any commentary on the POLES modelling or results, it only covers the way in which these results are used in subsequent analysis by DECC. This peer review was initially intended to provide a commentary on DECC's methodology used in 2011. However, it has become clear that more fundamental changes to the approach are needed to provide consistency between DECC estimates and current market prices. This note starts with a review of the previous methodology used in 2011. It then moves on to provide some suggestions on the use of market prices to inform the central scenario. Methodological improvements to DECCs model are then suggested in order to allow it to be used for policy 'what-if' analysis as the basis for developing high- and low-price scenarios for use in sensitivity analyses for policy appraisals.

1. Rationale and context for DECC's carbon price modelling approach

The DECC approach to valuing carbon for government policy appraisal in the traded sectors was changed in 2009. The previous carbon pricing methodology was based on estimating the social cost of carbon, effectively aiming for a 'fair value' estimate. This was changed, so that values for the traded sectors are now based on market value estimates. A key rationale for the change was that carbon should be treated as simply another economic resource cost for the traded sectors, so that any policy that raises or lowers emissions in those sectors will incur or save the corresponding costs. Since the EU-ETS market determines the price of carbon, policies should therefore evaluate carbon emissions according to market prices in the EU-ETS, irrespective of the UK government's view on the fair value of carbon. The pro's and con's to this approach are outside the scope of this review, which focuses on the approach to estimating market prices.

Forward prices in the EU-ETS go out a few years, but not as far out as 2020, so a model based on market fundamentals was developed as a basis for determining traded values. The 2011 model methodology was tailored to allow the near-term modelled carbon price values to match the actual market values.

Given that models can only ever be an approximation to reality, these kinds of pragmatic assumptions and model fixes are bound to be necessary to match models to real behaviour. The problem DECC faces is that real behaviour changes depending on market conditions, so the model structure used to assess prices in one year may be inappropriate the next, requiring regular methodology updates.

2. Review of 2011 Methodology

DECC's short-term price model uses marginal abatement cost curves (MACCs) provided by third parties as the basis for estimating the cost of achieving a given level of abatement in a given year. Given the ability to bank allowances between years in the EU-ETS, the annual MACCs need to be converted to cumulative to give an estimate of the cost of achieving an aggregate abatement effort over a particular time horizon (in this case 2011-2020). In the 2011 methodology, this aggregation was done by reading across all the annual MACCs in that time period at a particular price level, and adding up the emissions reductions in each year to give total abatement potential at that price. The total aggregate emission reduction required over the period to meet the EU-ETS cap was then compared to this aggregate abatement MACC to estimate the average price required over the period to achieve the necessary emissions reduction.

This average price was then taken to be the market price in the mid-point of the time period (2015). In order to derive price estimates for other years, assumptions were then made regarding the time profile of emissions in relation to the 2015 price, reviewed below.

More generally, this approach suffers from the problem that the accumulated effort derived from the MACCs assumes a constant price over the time period in question, whereas the final price profile is quite steeply increasing (£13/tCO_2 for 2011 rising to £29/tCO_2 for 2020). Given the non-linear nature of MACCs (e.g. effort in the longer-run is often cheaper than in the short-run because more options become available), assuming a flat price for all years is likely to give a rather poor estimate of the cost of achieving a particular level of emission reduction. It would be better to derive the aggregate emission reductions from the MACCs using the final carbon price profile, as this would ensure that the carbon prices across all years were consistent with the MACCs. This option is discussed further in the next section.

The price profile in the 2011 methodology is derived by assuming a cost-of-carry annual price multiplier, and an assumption about myopia amongst market participants.

Cost-of-carry. The first adjustment made to the carbon price trajectory is to assume that it follows the time profile of the carbon futures market. Given the no-arbitrage rule, futures prices are directly linked to spot prices by a nominal interest rate plus a cost of carry (i.e. the cost of borrowing money to buy allowances today so they can be sold next year). Given that allowances are bankable and free to store, this results in a very flat price trajectory in the futures market. The 2011 DECC model assumes a cost-of-carry of 1.5%, so carbon prices rise in real terms at only 1.5% per year over the period to 2020. This price profile is much too flat to reflect market risks.

The assumption that prices out to 2020 should follow the same no-arbitrage time profile as the futures market is flawed. Futures prices only go out a few years because the price risks become too high to sustain the liquidity required for the no-arbitrage rule to apply. A more useful pricing model over these timescales would be to ask what rate of return an investor buying allowances today would require in order for them to take the risk of a speculative long position in the market. This (real) rate of return could then be used as the basis for a risk adjusted discount rate that could be used to derive the time profile of carbon prices. Further thought should be given to appropriate discount rate to use.

Myopia. The 2011 model then assumes that the annual effort over the period 2011-2020 would deviate from this rather flat no-arbitrage price trajectory because of myopia amongst market participants regarding the level of abatement required in future periods. In other words, the market is assumed carry out less abatement in the current time period, and to defer abatement until later. The degree of myopia appears to have been calibrated to achieve a 2011 modelled price that was close to market prices at the time.

Given that the model was projecting a higher price level than the market was displaying in practice, this assumption of myopic behaviour is an understandable pragmatic step to 'correct' the model. It may have matched market commentaries at the time which were invoked to explain why market prices seemed low compared to fundamentals.

However, the myopia assumptions could probably be dropped. For the 2011 model, it seems likely that the model was producing a value that was inherently too high for the current period. The assumption of a flat carbon price over the whole time period would probably tend to overestimate the overall costs of achieving a particular aggregate level of emissions reductions. This is compounded by taking the very low cost-of-carry assumption instead of a larger risk-adjusted discount rate to calculate the time profile for carbon. Correcting these problems should obviate the need to invoke myopia.

In conclusion, the methodology for aggregating MACCs in the 2011 methodology was flawed, and should be corrected. Options for doing so are outlined in the next section. The assumption that carbon only incurs a minimal cost-of-carry should be replaced by an assumption that real carbon prices will rise at a full discount rate, and more thought should go into deciding what value of discount rate should be used. The assumption of myopia was a pragmatic step taken to help calibrate the model to market conditions at the time. As will be discussed below, this situation has changed radically, and the myopia assumptions could now probably be dropped, as other choices regarding calibration will need to be made.

3. Recommendations for 2012 methodology

The DECC model is geared up to measure the cost of abating emissions from their business-as-usual level to meet the cap. The model is not therefore well-equipped to deal with predicting prices in an over-supplied market. An alternative approach is therefore to use market prices directly as the basis for DECC short-run price forecasts. This option is discussed in Section 3.1. The DECC model will nevertheless be useful for undertaking policy 'what-if' analysis, and to develop high and low price scenarios. The changes required to the DECC model to improve the methodology are discussed in Section 3.2, and the choice of scenarios in Section 3.3.

3.1. Use of Market Prices

An alternative to trying to calibrate the DECC model to match current market prices would be to simply use current market prices as the basis of the short-term prices. One problem with this approach is that whilst futures prices are quoted out for several years, the volume of trades drops off significantly after 1-2 years. The information content of the traded prices is therefore poor over the longer time period to 2020 that DECC prices are required for.

One way round this is to extend prices out to 2020 from their current level using an assumed discount rate. The choice of discount rate is not without its problems, since unlike the short-run futures term structure, the long-run rate of return expected by the market is not an observable variable. The choice of discount rate has a significant bearing on the carbon price profile and the implied abatement schedule. Ideally DECC should complement its fundamentals analysis of supply and demand with a fundamentals analysis of risk in order to provide a bottom-up estimate of an appropriate risk adjusted discount rate. However, this would take some time to develop.

The rate of increase in prices for allowances of different vintage in the futures market does not provide an appropriate real discount rate on the future carbon price trajectory. Instead it represents a nominal rate which in theory represents the risk-free rate plus a cost of carry. The implied cost-of-carry has usually been small for carbon, since it is free to store and does not incur any dividend payments. The change in prices from one vintage to the next in the futures market has therefore historically been small, reflecting the no-arbitrage rule of liquid markets that effectively equalises spot prices in real terms with the futures price over a 1-2 year time scale.

Beyond this short time horizon, a shortcut to detailed risk modelling would be to take a variety of market forecasts of carbon price, and derive a risk-adjusted discount rate from the carbon price profiles contained in those. Alternatively, assumptions used in other areas of government appraisal regarding private sector discount rates in relevant sectors could be used as a proxy for the discount rate for carbon until a more fundamental analysis could be carried out.

In principle, the latest market price provides the best estimate of future price expectations. Basing prices on an average of historical prices over some time period to smooth out the effects of price volatility might make sense for commodities that display seasonal variability, but this is not the case for carbon. Such an approach is not consistent with the view that current market prices already contain all the information gathered up to that point in time regarding future expectations of prices.

One disadvantage of using market prices is that they will shift over time risking the DECC forecasts being out of date as soon as they are published. There could also be potentially large changes in the carbon price forecasts from one year to the next, which could be disruptive from a policy appraisal point of view.

Another down-side is the difficulty of generating different price scenarios. The simple approach of projecting forward from current prices at different discount rates to generate high and low price scenarios risks underestimating the range of future outcomes. For example, under this approach, prices would always increase from their current level. However, in reality, it is perfectly reasonable to consider scenarios in which the price might collapse further if the current oversupply of allowances continues (see Section 3.3).

For both these reasons, I would recommend DECC should therefore use their fundamentals model to develop high and low scenarios to be used as sensitivity cases for policy appraisal. The fundamentals model will then help put the market prices into context, and provide insights into the potential effects of different policy outcomes at the EU level.

3.2. Changes needed to DECC model

The model methodology should be altered to create more consistency between the aggregated effort derived from the annual MACCs and the final time profile assumed for carbon prices. There are two separate issues of consistency to consider:

A. MACCs are produced by external models, and show the level of abatement that can be achieved for any given price level in a particular year. These models assume some trajectory for carbon prices prior to the particular year that is being interrogated. The MACC will give the level of abatement achieved at a given carbon price, but there is a second-order error introduced if the carbon price trajectory being used to interrogate the MACCs does not have the same trajectory as that assumed in the model that was used to derive the MACCs in the first place.

B. The level of abatement achieved in a particular year is determined by reading across the MACC according to the relevant price. Clearly it is important to use the right price level in each year to get the appropriate level of abatement in each year. Errors in the price level will result in first-order errors in the level of abatement derived from the MACC.

This leads to two different options for improving the 2011 methodology. Given that option A would only correct a second-order error, and option B would correct a first-order error, I would recommend option B.

Option A. The MACCs this year are being produced by the POLES model, which generates MACCs by applying a linear increase in carbon price over time, and then assessing the abatement potential at particular point in time. The first option for improving the DECC model is to keep the methodology broadly the same as in 2011, but instead of using a flat carbon price over the whole 2012-2020 time period, to use a linearly increasing carbon price. This would help resolve the second-order error by improving consistency with the POLES methodology. The MACCs would still be aggregated to give total effort over the 2012-2020 period, and as was done in 2011, assumptions would then be applied regarding the time profile of carbon prices, based on cost-of-carry (and perhaps myopia).

Option B. In order to resolve the first-order error, the final carbon price trajectory should be used directly to read across from the MACCs the level of abatement available in each year of the 2012-2020 time period. Under this option, the shape of the carbon price trajectory (but not its actual level) would need to be determined first. For example, it might be assumed that prices follow a path shape determined by a risk-adjusted discount rate (see below for a brief discussion). Multiple price trajectories with this same shape would then be applied to the MACCs to determine their implied total emissions abatement over the 2012-2020 time period. The particular price trajectory that matches the level of effort required to meet the cap would then be selected as the 'correct' price level to be used as the output from the model.

3.3. Choice of 'What-if' Scenarios

The POLES baseline scenarios project emissions to be below the EU-ETS cap until the 2020s. Under the 2011 methodology which only considers the balance of supply and demand up to 2020, this would result in the DECC model projecting a zero short-term traded carbon price. Given that current prices in the market are low but not zero, there must be factors driving prices other than the current balance of supply and demand. The choice of modelling approach that DECC takes would depend on what they believe is driving price behaviour. Three different interpretations of the current non-zero carbon price, and corresponding modelling approaches are given below:

- Current positive prices may be a response to expectations of future scarcity in the 2020s. This could be investigated by simply extending the current methodology out to include MACCs up to 2030. In practice, given the high volume of banked allowances from Phase III implied by the POLES baselines, it seems likely that carbon prices arising from this approach would still be low given the weak aggregate demand for EUAs out to 2030 if the full amount of current oversupply is allowed to be banked into Phase IV.
- The current positive carbon price may reflect the option value associated with expectations of policy interventions such as raising the overall abatement target or withholding allowances to reduce oversupply. This could be modelled by running the model under different cap scenarios. Multiple scenarios could be combined and calibrated to current market prices to derive the implied probability weighting that market participants might be applying to such policy outcomes.
- The rational carbon price ought to be zero, but it is taking its time to get there. This delay in prices reaching zero may be a result of inertia in the price discovery process (such as happened towards the end of Phase I when prices gradually drifted down to zero), or could be linked to behavioural aspects of market participants. For example, industrial participants who tend to have excess allowances may be acting as compliance rather than market players, holding on to current allowances irrespective of future price expectations. Myopia could be invoked in this different context as an approach to model a gradual decay of carbon prices.

It is beyond the scope of this note to recommend an exact choice of scenario. However, it seems reasonable that a low-price scenario might be based on a continuation of the current situation regarding oversupply of allowances in the market (coupled with suppressed demand). It seems likely that such a scenario would produce projected prices remaining at or below their current level, and potentially falling close to zero for a time. A high price scenario could explore policy options for tightening the market, either by removing allowances or tightening caps post-2020. This could perhaps be coupled with a scenario of more rapid economic recovery which might boost demand.

3.4. Other Issues

Text in this section is based on email exchanges regarding additional questions from DECC subsequent to finalising the report.

Methodology for deriving upper and lower price bounds:

Using historical data on carbon price volatility seems to be an inappropriate basis for estimating upper and lower bounds on future prices. Carbon 'volatility' in the EU-ETS has tended to be low for periods of time, interspersed with jumps as significant new events arise (e.g. information about the recession etc. etc.). If you take historical volatility, you will essentially be measuring the effect of a series of historical events over whatever time period you choose, and assuming that these give you some view of changes that could occur in the future. This is somewhat arbitrary. If you multiply the std dev of these ranges by some figure going forward, this becomes even more arbitrary. I wouldn't know how to calibrate these multiplication factors, and I'm not convinced that past shocks are a good guide to future shocks. I would suggest instead that looking ahead and taking a view on the big uncertain variables (demand, energy prices, carbon cap) in your fundamentals model is a more valid approach to constructing scenarios.

Estimating upper and lower bounds on projected prices should be done with the fundamentals model, not using market values approach. The only variable you have to play with in the market values approach is the discount rate, and I don't think you will reflect the true level of uncertainty over future prices using different discount rates (especially on the low-price side, since using negative discount rate is hard to justify, so your scenarios will always show an increasing carbon price).

Choice of base year

If you are going to escalate prices at 7%, then it doesn't matter which year you use as a base year (i.e. 2013 prices escalated at 7% would approximately intersect the 2014-2016 prices anyway). As you've seen from my note, I don't agree with escalating the projected price at the futures price increase rate, since this is supposed to reflect a nominal risk free rate rather than a real risk adjusted rate. In practice however, the difference is not as high as it would have been if the futures price was increasing at its historically lower rate of around 5%.

Choice of averaging period for base price

I don't think you should average historical futures prices at all, since historical prices reflect the expectations players had in the past when they were more ignorant about the future than they are now (at least in theory!).

4. Summary

Overall, the following actions to improve the methodology should be considered:

- Base the central carbon price scenario on current market prices using a risk-adjusted discount rate to extend prices out to 2020
- Use the model to develop high and low price ranges and what-if scenarios
- Extend DECC model to include 2030 MACCs, and explore alternative cap scenarios as a way of calibrating to current prices
- Determine the shape of the carbon price profile based on appropriate risk-adjusted discount rate assumption
- Use this price profile shape to read across from the MACCs to match the aggregate effort over the relevant time period

Issues for future consideration:

- Investigate more thoroughly the link between market risk and discount rates to get a better handle on how to model the time profile of carbon prices

Consider integrating the model of abatement volumes and costs more directly rather than going through the intermediate step of using MACCs to avoid issues of price path inconsistency