

Assessment of LCP's Dynamic Dispatch Model for DECC¹

13 July 2011

Executive summary

DECC commissioned LCP to produce a model of electricity dispatch from GB power generators and of investment decisions in generation capacity in GB out to 2050 (the Dynamic Dispatch Model or DDM). As part of their quality assurance DECC has requested an assessment of LCP's DDM against DECC's specification.

DECC's main requirement is to model investor behavior and its response to policy interventions in order to guide the design of policy and to understand how investment and operation are likely to respond to fuel and carbon prices as well as the policy environment. The DDM can model investment behavior based on rational forecasts of market prices for up to 15 years ahead. Comparisons of the expected profit with that realized confirm the rationality or consistency of the investment decision. It can also model investment decisions based on pre-specified investor assumptions, thus allowing an exploration of the credibility of policies – if investors do not believe that carbon prices will rise they may make different decisions than if they are convinced that they will.

The DDM is an impressive model that meets DECC's original specification well. Further useful features can be added with modest extra programming. As with any complex model it will require skilled use and its results will need careful documentation. These should identify the key assumptions causing differences from the (also carefully documented) base cases. It is important that the DDM custodians and operators are themselves aware of these limitations and provide quality assurance to others who use the model before its outputs can be released. This will be facilitated by dedicated operations research expertise within DECC, supported by a management strategy for using the model and drawing on that expertise.

The results examined in this assessment are understandable and provide the level of detail to drill down and examine what drives the results and where the critical assumptions lie. It is important to recognise the need for intelligent risk framing of all results, and the importance of that framing in packaging results for DECC clients, given the flexibility of the model and the time period over which it forecasts.

¹ By Professors David Newbery and Daniel Ralph, University of Cambridge. They would like to thank LCP, who have been extremely helpful in answering questions and responding to suggestions.

Introduction

DECC commissioned LCP to produce a model of electricity dispatch from GB power generators and of investment decisions in generation capacity in GB out to 2050 (the Dynamic Dispatch Model or DDM). As part of their quality assurance DECC has requested this assessment of LCP's DDM against DECC's specification, an abbreviated summary of which follows:

The model is required to report GB outcomes for carbon intensity, security of supply and prices and incorporate the impact of current and possible future policy instruments. The model will be consistent with the modelling of electricity demand in Updated Energy and Emissions Projections

(http://www.decc.gov.uk/en/content/cms/about/ec_social_res/analytic_projs/en_emis_projs/en_emis_projs.aspx). The model's estimates and projections will be regarded by DECC's stakeholders as robust at the national level. The model will have a transparent and comprehensive front end that does not require programming expertise to run.

The Model should:

- Be a comprehensive fully integrated power market dispatch model covering the GB power market;
- Take into account investor decisions, and have these endogenous to the model.
- Allow explicit assumptions to be input such as:
 - fossil fuel prices,
 - inflation,
 - foreign exchange rates;
 - demand assumptions;
 - projected effects on the load duration curves due to policy and technological change
 - CO₂ prices;
 - generation plant capex and non-fuel opex;
 - plant efficiencies; plant availabilities;
 - plant closures and new build plans;
 - policy incentives (e.g. ROCs, LECs).
- Output metrics that the model would need to deliver:
 - Generation mix (by technology)
 - New Build Capacity (by technology)
 - Plant retirements (by technology)
 - Wholesale Prices (peak and annual average)
 - Capacity margins and de-rated capacity margins
 - Expected Energy Unserved
 - Probability of brown-outs
 - Plant margins

- CO₂ Emissions (by technology)
- Demand (annual/peak)
- Intermediate indicators - to aid interpretation and Quality Assurance of results:
 - Load duration curves if these are endogenous
 - Distribution of prices over the year
 - Merit order curves
 - Spark spreads, dark spreads
 - SRMC of different technologies
 - Expected IRR² of investment in different technologies

In addition there are a number of specific requirements that are discussed below. The consultants engaged to conduct the review had three tasks:

1. Review and comment on the proposed approach and methodology against DECC's technical specification;
2. Review and comment on functional requirements to ensure that the model is scoped to deliver DECC's policy objectives; and to
3. Review and comment the initial model against the agreed functional requirements to ensure the model is "fit for purpose". This review is a functional review and excludes review of mathematics, programming, usability, training and model documentation, and project management.

As part of the task of assuring that the model is "fit for purpose" and satisfies the agreed specification, the reviewers were asked to check that the results are explainable given the input data and assumptions and to comment on the limitations and strengths of the model. This would include advice on how the DDM should be used and handled within DECC once handed over and suggestions for possible extensions.

The review concludes that the model is a carefully constructed and flexible tool with a number of attractive features. DECC's main requirement is to model investor behavior and their possible response to policy interventions to guide the design of policy and to understand how investment and operation are likely to respond to fuel and carbon prices as well as the policy environment. This is a challenging task and the logical place to begin the peer review.

² IRR = internal rate of return, a measure of the profitability of the investment

Modelling investment

The critical element of the model is its ability to model investor behavior. DECC's requirements are demanding, and reproduced with numbering for convenience:

Investment decisions

1. Needs to distinguish between different stages of investment i.e. planning and then development.
2. Needs to capture the different market risks for different technologies
3. Needs the investment decision making process to be consistent with corporate finance theory and practice.

Uncertainty and Investor Expectations

4. Needs to be able to be run under perfect foresight and also without.
5. Needs to be able to capture investor uncertainty around central expectations for wholesale electricity prices, fossil fuel prices, carbon prices; load factors by technology e.g. wind.
6. Needs to capture how the uncertainty to returns to investments in different technologies varies under different policy regimes and the impact on investment in new generation.
7. Desirable that investor expectations for wholesale electricity prices should reflect the implications for capacity margins of demand projections, expected retirements, and new capacity under construction and planned as well as fuel costs and the policy framework. If for example the model is projecting increased volatility in wholesale electricity prices over time reflecting increasing renewable capacity then it would be desirable for investor expectations to anticipate this.
8. Desirable to capture varying investor views on the credibility of future policy and the extent to which policy commitments can be used to underpin credit ratings and bond finance.
9. Important that the modelling of uncertainty is transparent.

Requirements 1, 2, 5, 6 and 7 are handled appropriately in the DDM.

Requirement 3 is handled through user specification of the appropriate project discount rate which will need to take account of the portfolio exposure of the investor, or of the stock market assuming that a merchant investor issues shares.

This can be validated by studying the project profit profiles after the project run if necessary, by looking at correlations of project profit with total generation profit, or for the integrated firms, with their overall profit, perhaps taking a relatively short forward look consistent with the time horizon of the investment decision.

The way the model does this seems sensible, in that a competent modeller can make intelligent use of the input specification, and specifically the hurdle rates, which are computed as

Hurdle rate = Cost of Debt * (Gearing Ratio) + Cost of Equity * (1 - Gearing Ratio) + Technology Premium + Any policy hurdle rate adjustment, adjusted for duration of the plant's lifetime where the policy is applicable.

The hurdle rate can thus be varied by technology, and that may be enough to capture the technology-specific risks, although a good modeller will need to test the sensitivity of investment in each technology to these parameters to gain comfort. For example, the project risk of an open-cycle combustion turbine that can be built quickly and moved elsewhere if not needed is lower than a nuclear power station that could take 4-15 years to commission from ground breaking, and that can be reflected in the Technology Premium. Good practice should be to include footnotes or comments on why the hurdle rates have been chosen at the particular level.

Requirement 4 is the most challenging, and given the curse of dimensionality in a stochastic model, requires some intelligent short cuts. The model has been modified to look ahead for varying periods up to 15 years and simulate the prices that prevail in order to assess the profitability of the investment. The ability to compare expected and outcome IRRs gives reassurance about the accuracy of the perfect foresight assumption and in the examples tested the match is impressive.

There appears to be a rich range of possible assumptions that investors can be deemed to hold, including some like the future carbon price trajectory that relate to credibility of policy announcements. The model appears to have the desired capability to differentiate between perfect and investor expectations, and also be better placed to simulate the way in which other electricity market models with investor decisions handle investor expectations. At present investors can hold expectations about the future evolution of the carbon price, and the model can reflect their assumption that they do not believe the carbon price will rise, and the model does this by assuming that their expected future carbon price starts from the present level and remains flat. It is rather more difficult to model expectations in which they assume that the near term carbon price starts from where it is and even increases slightly over the next few years, but then the ETS and/or Kyoto process fall apart and the carbon price falls to zero thereafter. This can be accommodated by setting the low carbon price trajectory to follow this expected path, leaving the middle and high

carbon prices unchanged, and then concentrate on the lowest 5% of price outcomes, but this approach will need some careful handling to check that it captures investor expectations properly.

Requirements 8 and 9 are nice to have, and hence not a binding condition. Credibility might be addressed by assuming that investors ignore some future policies, which default in the investor's perception to the most adverse from the investor's viewpoint – e.g. without a carbon tax for zero-C plant and with a (possibly more modest?) carbon tax for fossil generation. Some form of ranking might be that any tax is the least credible, EU requirements moderately credible, contracts that are out of the money also moderately credible (and might thus be given a shorter life than it says on the tin) and contracts in the money most credible (longest duration). But I think DECC should take the lead on what they want here – perhaps they just make a larger number of runs with and without policies to see what difference they make and then they can make a judgement about their credibility separately.

Transparency of the DDM is probably as good as can be expected of any complex model.

Attractive features of the DDM

All models require simplifying choices, and they need to be chosen carefully if the model is to be fit for purpose. Aspects which provide such reassurance include:

- Taking account of correlations between coal, gas and carbon prices is reassuring;
- Treating demand side response as storage. This will require intelligent modelling and interpretation of the input specification;
- Modelling the wholesale price as short-run marginal cost (SRMC) + premium and considering a premium based on the tightness of available capacity. This will require intelligent modelling to set the right mark-up for each degree of capacity tightness. Note that the tightness is relative to actual capacity available, so that low wind hours put greater stresses on the price, as they should;
- Defining reserves as highest cost SRMC and top of the merit order, presumably this will evolve as some older plant moves up the merit order;
- Producing the SRMC including the carbon cost, as this clearly shows when unabated plant starts to undercut abated plant in the merit order.

Issues that will need careful modelling

The model is both sophisticated and flexible, but this can mislead users if they are not familiar with how it works and what its limitations are. It will be very important

that those who are custodians and operators of the model are themselves aware of these limitations and provide quality assurance to others who use the model before its outputs can be released (see further comments under Organisational Caveats to follow). To give a simple example, the RAB approach³ to supporting investment might appear to be the cheapest, but the standard regulatory incentive problem is that rate-of-return regulation does not encourage cost minimisation and leads to gold plating, so that the cost assumptions, which are data, should not necessarily be the same for different support instruments. One might wish, for example, to allow merchants to access cheaper construction costs if that is how they were able to enter the market.

The model has, understandably, solutions⁴ that are constrained by financial and physical limits to avoid silly answers. This is not too serious in that if an investment is not now profitable, there will be none done (in any one realisation) until prices rise as a result of delayed investment – sooner or later something will be built. But users need to be aware of how the model actually does this. Some of the questions to be explored, for example, the impact of demand side management (DSM), may be addressed rather indirectly by modelling them as storage options, and care will be needed to ensure that where the adaptation is a rough approximation this is clearly explained and caveated. At present one such issue is the treatment of interconnectors, which does not appear to be price sensitive. This is potentially important, as the ability to export rather than spilling wind and the ability to import in windless conditions may have a considerably impact on prices and the risk of unserved load.

Strengths of the model

- Clean modular structure means that it should be relatively easy to add new options and create new spreadsheets to view different aspects of the data. It would be helpful, e.g., to generate a cumulative cost path from the underlying data.
- New technologies such as peaking combustion turbines and different wind options (reflecting differing capacity factors from windy to less windy locations) can be added quite easily.

³ That is, paying a regulated return on the regulatory asset base (RAB) in the same way that other infrastructure such as transmission grids is remunerated.

⁴ Financial or physical constraints will limit the maximum rate and utilities and merchants may face different constraints.

- Price determination appears to be sound, as noted above. This is consistent with supply function models; but might be problematic if there is no high SRMC plant actually on the system (although given the need for flexible reserve gas plant plus a carbon cost this is probably not an issue).
- The impact of build constraints, which are very powerful for profitable new generation, can be clearly seen in the capacity evolution of each technology. (This is best seen by running the stochastic version in single run mode by turning off the stochastic features.) This should allow adjustments to be made to policy support instruments so that investors do not receive excessive wind-fall gains – here the graph of IRRs is helpful, as will be adjusting the merchant hurdle rates to simulate slightly more pessimistic financing requirements to see when vertically integrated utilities (VIU's) invest but merchants do not.
- The ability to compare the IRRs that guide investment and the out-turns (although not on the same graphs yet) is helpful. The graph of IRRs gives an indication of when technologies become commercially attractive. It might help if the technology specific hurdle rates are graphed on the same figure.
- The clear documenting of assumptions is good. This will enable DECC to ensure that different models and projections share a common periodically updated set of assumptions. (It would be helpful here if the policy assumptions could also be viewed in the output file – for example, the carbon prices are in the outputs but they are not the complete story as the carbon price floor (CPF) will be binding in some periods. Thus the carbon price page might have three lines: the ETS price, the CPF uplift (i.e. the extra carbon tax) and the resulting CPF, which may be the same at the ETS price.)
- The graphical display of the percentiles of the stochastic outcomes is clear and should allow other decision criteria (minimum regret) to be examined. For example, banks may be interested in projects where there is a less than 5 or 10% chance the IRR will fall below some level. The hurdle rate can be adjusted to deliver this.
- The use of test technologies to sample their attraction over time is clear and helpful.
- The ability to compare different runs in split screen mode is good and the exported spread sheets should allow a variety of comparisons to be generated.

- As an important mutual validation exercise it seems fairly straightforward to input the same assumptions as other models and compare outputs. The backcasting test discussed in the appendix is a good example.
- It appears that the investment behaviour modelled allows for some investigation of credibility in the following way. If “perfect foresight” is turned off for the carbon price assumption (as the most likely doubt would be around the future carbon price) then this can be set either to grow at 0% or some lower rate from the date of investment, and the impact compared with the full “perfect foresight” assumptions, and also by comparing the expected and outcome IRRs. This demonstrates a sharp difference between fixed FiTs or CfDs,⁵ where the future electricity price is less relevant, and ROCs cum premium FiTs⁶ where it matters and depends on the carbon price (at least at modest carbon prices where the choice is sensitive to the carbon price).
- Plant build cost reductions do not respond to cumulative capacity but are pre-determined. This is a good feature in forcing the modellers to revisit and cost assumptions in the light of investment history – plant that looks uneconomic perhaps should not enjoy rapid cost reductions if no-one wants to build any.

Organisation caveats - conditions of use

1. DDM team. Managing and running the model primarily within a dedicated DECC group, the “DDM team”, has the advantage of providing due diligence based on an intimate understanding of the assumptions built into the model. This will, undoubtedly, be greatly facilitated by dedicated operations research expertise within the DDM team supported by a management strategy for making use of this expertise and the model within DECC.
2. Quality assurance. At the risk of reducing either the availability of the model within DECC or the speed of response of the DDM team to DECC clientele, quality assurance of model outputs is a responsibility that naturally falls to the DDM team. A process for meeting a client brief should be developed and documented and potentially revised over time. For example, this process could be based on three stages.

⁵ FiT is a feed-in tariff, normally interpreted as offering a fixed price for possibly variable output, while a CfD is a contract for difference, paying the difference between the strike and market price.

⁶ ROC – Renewable Obligation Certificate paying a premium above the market price, like a premium FiT or pFiT

Stage 1: A short initial consultation results in a document to briefly state the question posed, assess the appropriateness of the model to address it, and note the main limitations of the output reflecting the simplifying assumptions of the model. In some cases, this assessment might conclude that the outputs requested would need to be checked or substantiated by a parallel or secondary process, perhaps outside the DDM team.

Stage 2: The DDM team sets inputs and runs the model including what might be called risk framing via sensitivity and stress testing. Where the brief is to propose or test viability of meeting a policy target, the goal of risk framing may be to identify circumstances in which there is a failure to meet the target.

Risk framing could include

- i. Sensitivity of key model outputs, such as targets, to input data.
- ii. Trialling the model in several extreme scenarios, aka, worst cases.
- iii. Looking at stochastic runs and adjusting the policy mix or level to meet the desired targets with 50%, 75%, 95% chance. Differences between model outputs should be traced back to those of model inputs.

Stage 3: A consultation with the client to present a response to the brief.

Beyond presenting the results and the qualifications, from Stage 2, the purpose of Stage 3 is to allow another iteration of modelling. Depending on the client, overall progress may be faster with several quick and dirty iterations between stages 2 and 3 to identify the best match between the model's capabilities and the client's needs.

3. Base cases. Having more than one base case, and informally surveying or testing a client's understanding of what each represents, may be helpful in providing quality assurance. Initial thoughts include a stagnant carbon price and/or a stagnant gas price. These could be combined with policy settings that ensure a 20% cut in greenhouse gas emissions and/or a 20% increase in the use of renewable energy by 2020. This suggestion may be developed in the first few months of operation, prompted by validation of model outputs.
4. Dissemination with DECC. The DDM team could give some thought, over the first few months of operation, on to how to present, and prospectively re-use, results from runs of the model. The starting point is intuition (e.g., in the base model, continual growth of carbon prices seems to fix the carbon emissions problem by 2050 irrespective of policy measures, and that fixes the wholesale prices, which are insensitive to carbon and fuel prices given the high degree of decarbonisation.) Next comes reports written for clients – easy enough to

make these available within DECC, and a good instructional tool where the report could give explanations about any counter-intuitive findings. More demanding would be a way to index into past runs rather than past reports, though it may not be worth effort to set up a software system for this. A much easier step would be to allow prospective DECC clients to search on difference between input files, used in past analysis, and the base cases. (The 2050 Pathways model provides examples of this.)

5. External dissemination and review. For example, the Stanford Energy Modelling Forum <http://emf.stanford.edu/> compares various long-term projections generated by different modelling groups using a standard set of input assumptions.⁷

Other caveats

The demand forecasts and their elasticities are an input assumption to the model and have not been explored here, but clearly there may be important feedbacks from retail electricity prices to aggregate demand that should be modelled (and there is some evidence that the long-term price elasticity may be quite high, 0.5-0.8 in absolute terms). Clearly demand reduction is part of the climate change policy and is at least some compensation for higher prices, and will be important for determining consumer bills.

Issues for possible model extension and desirable additions

There are some possibly relatively simple additions that might be helpful:

1. An index of parameters, inputs and outputs and their location.
2. Units clearly specified in all tables (e.g. TWh/yr or /qr, GW, bcm/yr, etc)
3. Cumulative investment cost
4. NPV of operation costs and investment costs of particular policy options at a predetermined discount rate, to allow comparisons across scenarios – e.g. the

⁷ “EMF 24 will focus on development and cross model comparison of a new generation of comprehensive international and domestic climate policy intervention scenarios focusing on technology strategies for achieving climate policy objectives. These scenarios will enable the community to exercise enhanced modelling capabilities that were focused on in previous EMF studies on the international trade implications of climate policies, the representation of technological change, and the incorporations of multi-gas mitigation and land use emissions and mitigation policy alternatives. ...Attention will be devoted to analyzing the efficacy of different policy architectures including different assumptions about the timing of mitigation, the allocation of mitigation obligations across nations, and the choice between tax, economy- wide and sectoral emission caps, technology standards, and GHG intensity standard policy instruments.”

cost of not having any nuclear. The DECC DDM team has the ability to write suitable programs to extract this information (as well as the previous suggestion).

5. A simple way to check what assumptions have been changed. If it were possible that all the assumptions were gathered on one spreadsheet then simply looking at whether cells have been changed might be enough. This will be particularly important where more than one person or group is using the model to try out options and alternatives, in documenting what input changes caused what output changes. Again the DDM team seem confident they can write the appropriate visual basic macros.
6. The ability to compare on the same sheet the expected IRRs and the realised IRRs – either on a summary or by technology. At present it is hard to locate individual plants (technology and date) from the long list, but it should be possible to insert filters that simplify finding and comparing across projects.
7. The price determination for tight supply-demand balances will be important for determining the amount of peak capacity built and this will depend on good modelling of low wind conditions, as the main stress may arise when wind penetration is high but its output low, and when some other plant has had an outage, not necessarily at system peak. The main issue here is to ensure that the modelling of the daily wind load duration curve has enough segments to include some near zero outputs even if they are rare.
8. Exploring the scope to automate sensitivity analysis? Varying “one parameter at a time” is the current starting point and could be done to produce tornado charts although this may be time intensive in the context of stochastic model runs. Given its time demands, this might be done relatively rarely, e.g., as part of a final quality assurance check before certifying model outputs. Using Monte Carlo Simulation to see the effect, graphed in percentiles, of simultaneous parameter variations is another possible approach although would require additional programming and testing.

There are more substantial changes that probably require significant programming and it will be for DECC to decide whether it is worth elaborating the model or finding fixes to substitute.

1. At present it is not possible to examine annual or daily price duration curves for the same generation park but different wind scenarios. This was not

included within the original model specification, but DECC may wish to commission LCP to add suitable additional functionality.

2. The model is a copper plate and cannot model grid congestion, although perhaps two versions (England and Scotland) could be run iteratively or even interactively. It may be better to use the model to gain some indication of total wind capacity (of differing capacity factors) and then test these against resource and grid constraints explicitly as these are likely to be site specific to be worth modelling in detail.

There are a number of off-line calculations that it might be possible to do if the output files are easy to access. One might be to provide an evolutionary view of aggregate VIU balance sheets and even Profit & Loss, given investment, gearing, profits and the cost of debt, all of which are either inputs or outputs. Depreciation, taxes, borrowing, new equity could then be separately examined. This would not have to be built into the model, merely some indication of how to pull out the elements needed.

In deciding for what purposes the model is suited, these various strengths and limitations need to be kept clearly in mind; and again the DECC modelling team should ensure that the model is not used inappropriately for questions for which it is not well-suited.

Possible questions for which the model seems well suited

1. Exploring the way in which the build rate and low-C plant mix might respond to a combination of incentives given various scenarios for fuel prices and construction costs, and the rate at which carbon targets are met and at what cost (complements the electricity model of the 2050 Pathways model).
2. Testing the ways in which the prices of FITs and CfDs can be adjusted to reduce the cost of support, given varying assumptions about investor behaviour and financial structure. (This will need a skilled modeller, given the abstractions incorporated in the structure of the model - e.g. in practice investors will all differ in their expectations, access to finance etc).
3. Exploring the impact of wind on security of supply given varying amounts of reserve capacity, including interconnector capacity (although modelling the use of interconnectors will take some care as at present it is non-price responsive).

4. Exploring the displacement effect on private reserve capacity build of publicly tendered reserves.

Possible questions for which the model seems less well suited

1. Determining the amount, location and value of extra transmission to relieve congestion caused by new plant (but National Grid is developing such a model and this is almost certainly something that should be done with a very different set of location specific data).
2. How much investment in interconnectors might be driven by wind, the carbon price floor (CPF) and other policies to stimulate low-C generation investment – again this might best be done off-line as the value of interconnectors depends sensitively on the price characteristics over short periods of time at each end – not so much on the average prices over a year or season.
3. Determining the value of reserves and ancillary services and hence deciding on their optimal level. Again these depend on the fine details of pricing, constraints, interconnection, and flexibility of plant.
4. The modelling of merchants will require some care in iterating between the capital cost assumptions and the risk implied by the form of contracting. At present the VIUs always pre-empt any investment because they have a lower cost of capital, at least until they are cash constrained. Debt and equity are inputs that can be varied and will jointly determine the weighted average cost of capital or WACC. These should reflect the form of contracting – again evidence that the model should be run by those familiar with the way it works, and will need careful validation if others use it to argue their case.

Conclusions

The DDM is an impressive model that meets DECC's original specification well and appears to have the capability to have some additional and useful features included with modest extra programming. As with all complex models it will require skilled use and results will need careful documentation that identifies the key assumptions that causes differences from the (very carefully documented) base case, or perhaps base cases. As far as the reviewers can tell it produces results that are understandable and provides the level of detail to drill down and examine what drives the results and where the critical assumptions lie. It is important to recognise the need for intelligent risk framing of all results, and the importance of that framing

in packaging results for DECC clients, given the flexibility of the model and the time period over which it forecasts. It is suggested that the DDM team aim to produce internal working papers setting out the results of varying assumptions together with an explanation of what is driving the results, so that expertise elsewhere in DECC can be mobilised to comment on and help improve the assumptions and policy mixes to be explored.