Is there still merit in the merit order stack?

The impact of dynamic constraints on optimal plant mix

Iain Staffell[1] and Richard Green

Imperial College Business School, London, SW7 2AZ, UK.

Abstract

The merit order stack is used to tackle a wide variety of problems involving electricity dispatch: assessing the impact of new technologies on prices and carbon emissions, or looking at the optimal investment in new capacity. The simplification it relies upon is to neglect dynamic constraints such as start-up penalties, energy storage, minimum stable loads, reserve requirements and maximum ramp rates.

We demonstrate a non-linear optimised dispatch algorithm and compare it to the merit order stack approach, using the GB system in 2020 as an example. We find that two constraints (start-up costs and minimum output levels) are the greatest source of error: the merit order stack under-estimates the optimal level of investment in both peaking plant and inflexible baseload generators, and thus their run-times by up to 30%.

A simple heuristic can be applied to the data on plant costs to account for start-ups, which reduces this error by a factor of two. This enables the benefits of the merit order stack (speed, simplicity and transparency) to hold in more dynamic and intermittent simulations.

Introduction

The UK's plans for decarbonisation assume that by 2030, a very high proportion of electricity will come from low-carbon generators, and that in the following decades, electricity will be used to meet a high proportion of the demand for transport and heat. The need to balance these much higher demands against the generation from inflexible sources raises many technical issues, and comprehensive engineering models are needed to address them. For example, Strbac *et al.* (2012) model the electricity system in detail to ask how much storage capacity might be needed by 2030, and whether it should be placed closer to generators (saving on transmission capacity) or to loads (saving on distribution capacity). In a similar vein, deciding whether an interconnector from Norway should connect to the UK system in Scotland or in Lincolnshire, for example, requires extensive modelling of the flows on the transmission system.

At the same time, the shift to a low-carbon energy system raises many economic issues. What are the consequences (for prices and for investment incentives) of adding 30 GW of wind to a system with a peak demand of 60 GW? What value should we place on increased interconnection with our neighbours? How will capital-intensive new build fare in a more volatile marketplace? Can a market that pays only for energy deliver sufficient peaking plant? Many of these questions depend on uncertain variables such as fossil fuel prices. Sometimes, a complex model can be solved for a few representative scenarios, but finding the distribution of a peaking station's profits requires a

[1] Corresponding author: i.staffell@imperial.ac.uk

model to be run many times. To avoid the computational burden of repeatedly running a full engineering model, economists frequently use simplified approaches. Finding an intuitive explanation for the link between inputs and outputs is often also easier with a simplified model.

One of the economist's workhorse models for linking generators' capacity and costs with their outputs and electricity prices is the so-called merit order stack. This ranks power stations in order of increasing variable cost, and always selects the cheapest available stations to meet demand. Investment decisions are made to minimise the sum of variable and fixed costs. This approach ignores the growing role that dynamic constraints will have on determining the optimal capacity mix. Plant start-ups and shut-downs, limits on ramp-rates and energy storage are three issues that cannot be incorporated into a simple merit order stack, and require a dynamic dispatch model with high temporal resolution.

To test the impact of these dynamic constraints, we demonstrate a non-linear dispatch optimiser coupled with a long-run investment model. This finds the operating pattern that gives the minimum annual cost of generation, and the mix of plants such that each type covers its costs, leaving no further opportunities for profitable entry to the market. A representation of the 2020 GB electricity system is used to explore the capabilities of this model relative to a merit order stack. Using these models, we ask how far the various non-linear constraints change the operating patterns for different plant, and the extent to which this would change the optimal plant mix.

Furthermore, we propose a relatively simple extension to the merit order stack model, modifying the input cost parameters to account for the expected number of plant starts per year. The impact of this extension on results is quantified, and limitations to its applicability are explored.

Background

The electricity industry has a long history of using sophisticated models to plan its investment and operating decisions. The Central Electricity Generating Board in England and Wales, for example, used one model to plan the operation of its power stations, contingent on their fuel costs, and iterated these results against a model of the coal industry and its transport costs to allocate the coal from particular mines to each power station (MMC, 1981). In operational timescales, when plant are being dispatched, it is vital to take account of constraints such as the time required for a station to start and limits on the speed with which it can change its output, and the CEGB used a computer programme known as GOAL (Generator Ordering And Loading) to do so.

For decisions made over longer timescales, these operating constraints may be less important, and simpler models can be appropriate. In particular, the problem of minimising the sum of investment and operating costs can be seen as the solution to a linear program, in which those costs are the objective function and the constraints require that the output from every power station is always less than its own capacity, and collectively equal to demand in each period. For such a model, known as a generator stack model, the stations with the lowest variable costs are used most intensively, running whenever demand exceeds the capacity of stations with equal or lower variable costs. Investment decisions are made by trading off fixed and variable costs, choosing plants with low variable costs when long running hours are expected, and those with low fixed costs if they are only needed to meet peak demands for short periods.

The underlying paradigm is that the sensible candidate plants for investment have the same ordering, whether ranked in order of increasing fixed costs or decreasing variable costs. A station with higher fixed costs, compared to some alternative, would only be a worthwhile investment if these were offset by lower variable costs and if the station was expected to operate for long enough for this to bring its total costs below those of the other option. A screening curve – the upper panel of Figure 1 – shows how the total costs (per kW of capacity per year) for three investment candidates vary with the amount of power they produce. The slope of each line shows the variable costs of that plant type, while its fixed costs are given by the vertical intercept. Peaking plants are cheapest if only a few operating hours are needed, while baseload plants have the lowest costs if power is needed throughout the year. The positions of the lines in this figure are illustrative; more details of the model are given in Stoft (2002) or Kirschen and Strbac (2004).

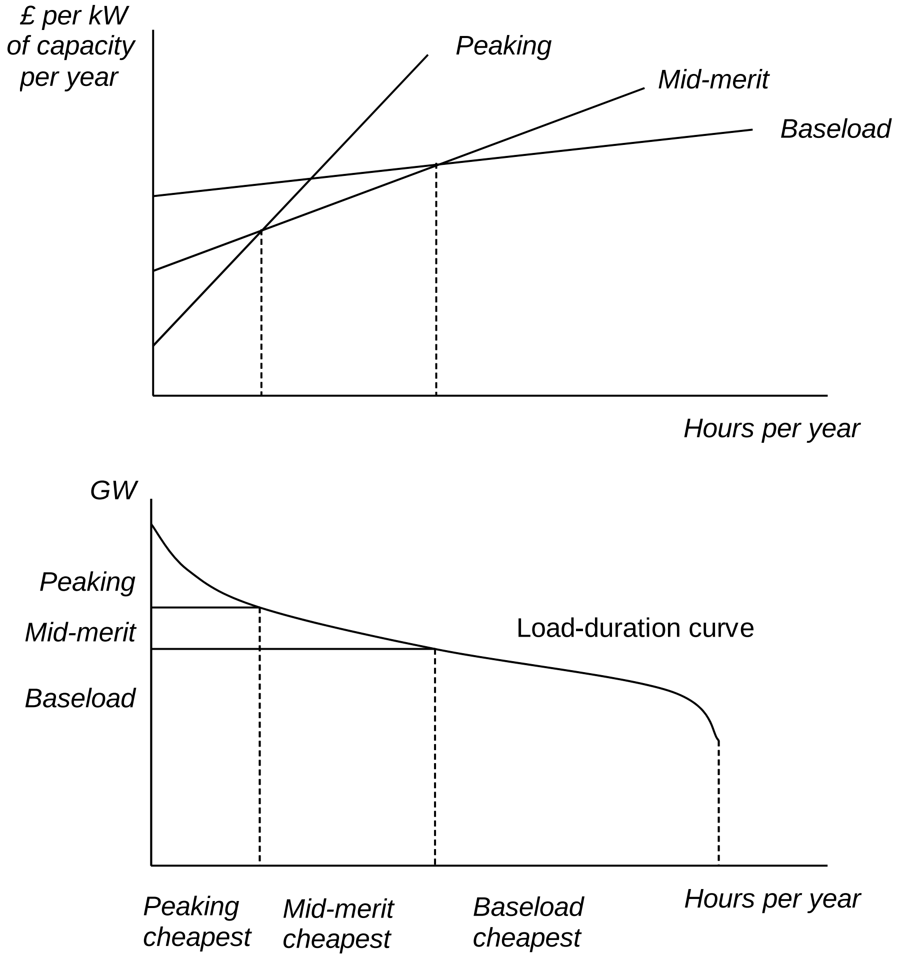

Figure 1: Merit order stack model (stylised)

The lower panel of Figure 1 shows how the model selects generator capacities. The load-duration curve ranks the levels of demand in each hour of the year, from highest to lowest. The chosen capacities of each plant type are stacked up against the load-duration curve, placing those with the lowest variable costs (the baseload plants) at the bottom. The more stations of a given type that are built, the fewer the hours that the least intensively used of them will be needed in. With the correct mix of stations, this number of hours is the number for which that type of station has the same total costs (fixed and variable) as the type with the next highest variable costs. Deviating from this mix of capacity would increase the total costs. For example, building an additional baseload station would mean replacing a mid-merit station that would have lower total costs for its operating pattern; building one fewer would imply that a mid-merit station was required to run for a number of hours where a baseload plant would have been cheaper.

The model described so far would require sufficient generating capacity to meet even the highest demand, even though some of that capacity would be used for only a few hours a year, which is unlikely to be economic. Traditional models allowed the demand constraint to be violated in a (very) small number of hours per year, equivalent to the reliability standard imposed on electric utilities. As an alternative, an option to reduce peak demands by shedding load can be introduced, as a kind of generation with a variable cost equal to the Value of Lost Load (VOLL). This cost should be set to a level which makes its widespread use unattractive. A third variant is to introduce price-sensitive demand, either through specific blocks of demand reduction that can be called off at particular prices or by making the level of demand in each hour a function of the market price. That market price is then set to the variable cost of the most expensive plant in operation or to the level needed to ration demand to the generation available at that price.

Merit order stack models have often been used by economists to simulate electricity prices in this way. For example, Joskow and Kahn (2002) and Borenstein *et al.* (2002) estimated the level of prices during the California electricity crisis assuming competitive (price-taking) behaviour and compared them with the (much higher) out-turn. Borenstein *et al.* took advantage of the relative simplicity of the merit order stack approach to run their model 100 times with different random plant outages, thus capturing the non-linear relationship between available capacity and prices. The merit order stack approach to costs also underlies the supply function approach used by Green and Newbery (1992) to study market power, and by Green and Vasilakos (2010) to model the impact of wind power on wholesale prices.

Merit order stack models can also be used to study questions related to investment. Borenstein (2005) uses a model of this kind to assess the benefits of introducing real-time electricity pricing in California, calculating the reduction in overall capacity and the change in wholesale price patterns that this would bring about. A number of papers have studied the impact of the large-scale introduction of renewable generators on investment in conventional plants – these include Lamont (2008), Usaola *et al.* (2009) and Bushnell (2010).

Mills and Wiser (2012) go beyond the merit order stack when they calculate the marginal value of wind power and of different kinds of solar power (with and without energy storage) at different penetration levels, calibrated for a scenario of California in 2030. Their model combines a long-term investment equilibrium and a short-term dispatch that takes account of operational constraints. Similarly, Strbac *et al.* (2012) calculate the value and optimal capacity of electricity storage for the UK in a low carbon 2030 scenario, using a simultaneous optimisation of investment and operation decisions, subject to plant- and network-level constraints.

These engineering approaches take account of many more constraints than the merit order stack, but involve much more complexity, making them less suitable for large-scale repeated simulations. We wish to test the impact of this complexity on the results obtained from a model of Great Britain in the 2020s.

Optimised Dispatch Model

We demonstrate a generic, non-linear optimised dispatch model written in the GAMS language and controlled by a web interface. The model consists of several classes of power station that are dispatched to meet a set of time-varying demands so as to minimise the cost of generation. Equally, the model could be used to maximise profits (simulating an unregulated monopolist), maximise welfare (benevolent planner) or minimise carbon emissions (environmentalist). The model copes with relatively large problems, optimising a fleet of 20 plant types over 8,760 hours of demand in around 20 minutes on a standard workstation (3 GHz, 4 GB RAM). Monte Carlo trials and multi-dimensional sensitivity studies can be conducted using the web interface.

Once plant characteristics, demand data and economic information are supplied, the model optimises the scheduling and output of plants subject to the following constraints:

- Demand, plus a reserve margin, must be served by the operating plants; failure to do so incurs penalties due to curtailing supply (spilling wind) or demand (charged at the value of lost load (VOLL));
- Price sensitive demand is modelled with several tranches of consumers (usually large industrial users) that are able to reduce load in return for a scarcity price (which is above the industry's typical marginal cost but below VOLL);
- Plants have a minimum stable output below which they must shut down. Restarting the plant incurs a cost and time penalty;
- Plants have minimum uptime and downtimes, and maximum rates of change in output;
- Plants have reduced efficiency when operating part loaded.
- Hydro and pumped storage are subject to availability constraints due to water levels;

The model finds the short-run equilibrium – how to best operate a given set of plants so as to minimise the total cost of generation. Long-run equilibrium – the capacity of plants that would be best to build – is found when the profits of each type of plant are closest to zero, and so there is no incentive for new capacity to open or for existing capacity to retire.

GAMS is not capable of performing nested optimisations (a model within a model), so an additional layer is required to find the long-run equilibrium. The web interface is used to conduct a simple iterative search: testing a set of plants, refining the levels of capacity based on their profits, then testing the new set of plants – until convergence is achieved. When considering a brown-field scenario (where existing plant operates alongside potential new plant) there is the constraint that the level of old plant (which has already been paid for) cannot be increased.

This model is similar to those demonstrated by Mills and Wiser (2012) and Strbac *et al.* (2012), and is available from the authors under the Creative Commons licence.

Data for UK Electricity in 2020

The model relies on five sets of data: the capacities, technical limitations and costs of each type of generator, plus time series of demand and output from wind generators.

Plant Data

We consider four types of renewable and seven types of controllable thermal generation, which are listed in Table 1 in order of merit (increasing marginal cost). The unit and total capacities are initially calibrated to the GB electricity system in 2010, except wind capacity which we project will increase from 7 GW in 2012 to 30 GW in 2020, in line with projections by RenewableUK (Green and Staffell, 2012).

Thermal plants operate with an availability of 85%, which is evenly distributed over the whole year. Wind output is modelled explicitly using a profile of resource availability, while hydro output is optimised, subject to the constraint of water availability, giving annual load factors of 42% and 15% for run-of-river and pumped hydro respectively.

Based on data from US generators (Bushnell and Wolfram, 2005), we assume that plant efficiency scales linearly with output, falling 6% from full to minimum output (for coal, OCGT and oil) or 16% (for CCGT and nuclear). When running partly loaded, the average generating efficiency of thermal plants are therefore lower than given in Table 2; however, the incremental efficiency (increase in power output from an increase in fuel input) is higher.

	Unit Capacity (MW)	Installed Capacity (MW)	Average Net Efficiency (%)	Minimum Stable Generation	Cold Start Time (hours)
Wind (onshore)	–	11,000	–	0%	–
Wind (offshore)	–	19,000	–	0%	–
Nuclear	500	11,000	35.1%	75%	96
Coal (large)	525	25,725	40.1%	50%	4
Coal (small)	250	3,250	36.7%	50%	3.5
CCGT (large)	750	23,250	56.9%	50%	2
CCGT (small)	350	8,400	51.2%	50%	1.5
OCGT	30	2,340	32.3%	10%	0.1
Oil	50	4,000	30.8%	10%	0.1
Hydro	–	1,360	–	0%	–
Pumped Storage	–	2,828	77.4%	0%	–

Table 1: Technical parameters for the power stations in our model

Cost Data

Plant costs were derived from five major studies: Mott MacDonald (2010), Parsons Brinckerhoff (2011) and Arup (2011) specific to the UK; plus IEA (2010) and EIA (2010) internationally. We aggregated their projections to 2020 or thereabouts for annualised capital investment cost (defined as the annual rent required to cover overnight capital cost plus interest over the lifetime of the plant), fixed and variable operating costs, and thermal efficiencies. Fuel costs are based on DECC's central scenario for 2020, which equated to £7.94 for coal, £37.07 for oil, and £23.37 for gas (per MWh of fuel). Carbon emissions are priced at £30 per tonne, which is the floor price established for 2020 under the government's carbon price support scheme (HM Treasury, 2011).

No-load costs are derived from the intercept of total fuel cost against plant output, and represent the penalty of decreasing part-load efficiency. Start-up costs are derived from the cost of fuel required to heat the generator to temperature plus the cost of the carbon emitted. The wear and tear caused by start-ups is not factored in; however, this could increase the start-up cost significantly (Rosnes, 2008). Shut-downs are considered to incur zero cost.

	Annualised Capital Cost (£/kW-year)	No-load Cost (£/h)	Start-up Cost (£/plant)	Incremental Fuel + Carbon Cost (£/MWh)	Total Fixed Cost (£/kW-year)	Total Variable Cost (£/MWh)
Nuclear	401	320	4,000,000	3.35 + 0.00	470	5.00
Coal (large)	207	620	142,500	18.62 + 25.84	240	47.78
Coal (small)	245	320	72,000	20.34 + 28.24	278	52.02
CCGT (large)	85	4,900	32,500	34.52 + 10.77	103	53.25
CCGT (small)	107	2,550	11,000	38.39 + 11.97	125	59.06
OCGT	58	130	270	68.12 + 18.99	72	93.30
Oil	146	360	470	113.07 + 27.23	188	149.27

Table 2: Economic parameters for power stations in our model

The last two columns of Table 2 summarise our cost data: the fixed cost consists of the annualised capital cost and fixed operation and maintenance (O&M) costs; marginal cost consists of fuel, carbon and variable O&M.

For this set of simulations we searched for a so-called “greenfield” solution; a long-run equilibrium assuming that there was no existing plant. The model can also accommodate a brownfield scenario when existing plant only has to cover its operating and maintenance costs, on the basis that the costs of past investments are already sunk.

Demand Data

Our hourly time-series of demand data was produced from half-hourly figures published by the National Grid for the period 1994–2011 (National Grid, 2011). We scale up this historic demand to hypothetical 2020 levels, assuming that demand will grow by 0.7% annually to give a total of 350 TWh per year. Historic demands are increased by the ratio of 350 TWh to their annual weather-corrected demand; in other words, we preserve hour-to-hour and year-to-year variation due to weather, while removing fluctuations due to the level of economic activity. The linear scale factors that we use do not reflect changing patterns in the underlying demand due to de-industrialisation and the potential electrification of heating and transport demands, so the system peak and minimum demands therefore also scale linearly to 64.6 and 22.7 GW respectively.

Wind Resource

The time-series output of wind generators was estimated using the methodology presented in Green and Vasilakos (2010) and Green and Staffell (2012). We obtained hourly observations of wind speed from the UK Meteorological Service (2006), collected from 120 weather stations between 1994 and 2011. Around 3% of the observations were missing or corrupt, and were filled using interpolation, regression and Markov-chain simulation.

Fleets of wind farms composed of 16 leading turbine models were stochastically allocated to each region, and the power curves from these turbines were mapped onto wind speeds to give the expected energy yield and load factor for each region. Wind speeds were adjusted to account for the fact that wind farms and weather stations have different hub heights, and reduced by 10% for onshore regions to give a mean load factor of 26% – matching the historic output of UK turbines. Offshore wind speeds were provisionally inferred from coastal locations, as none of the MIDAS stations were deployed at sea, and so were increased by 10% to give an average load factor of 36%. We expect that new speed measurements from Round 3 offshore sites will soon be made available, and will incorporate them into our data set. The diurnal variation in wind output was reduced to match historical observations, as detailed in Green *et al.*, (2011).

The final data set contained hourly load factors for GB, which closely resemble the pattern and distribution of actual output from transmission connected turbines in the UK. Figure 2 compares the spread in our estimated load factors with historical output derived from Elexon data and Renewables Obligation Certificate (ROC) submissions.

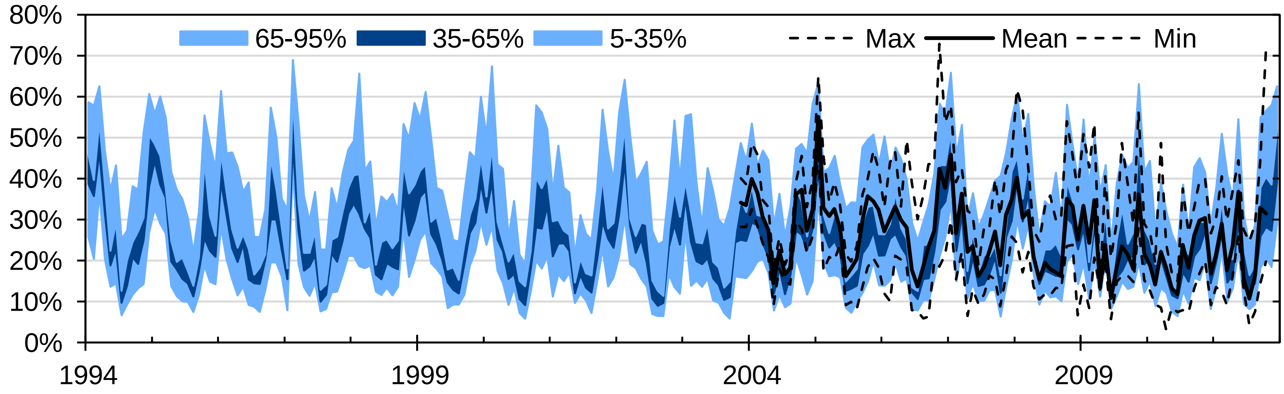

Figure 2: The distribution of monthly average load factor across the 120 wind sites, comparing simulation (shaded areas) with measured historical output (lines).

Extending the Merit Order Stack

The time required to run a full dispatch model on a year of data makes it unsuitable for many economic applications. One alternative is to use the dispatch model on a representative sample of days; Green *et al.* (2011) show how clustering techniques can be used to create these. Another is to use a full dispatch model to obtain more information on how the stations at each place in the merit order will be used, work out what this means for their costs, and then adjust the screening curve accordingly (Batlle and Rodilla, 2012). We propose a simpler method of extending the merit order stack to consider the cost of plant starts and the scheduling of hydro resources.

Plant Start Costs

We take a simple approach to estimate the number of start-ups that each plant will undergo: any time national load rises above a certain value, *Q*, we assume the *Q*th plant must start up. Figure 3 demonstrates this simple algorithm, showing how many times a plant must start at different positions within the stack. Using the UK as an example, the number of starts per year is highest for mid-merit plant which must run for around 3,000 or 5,500 hours per year.

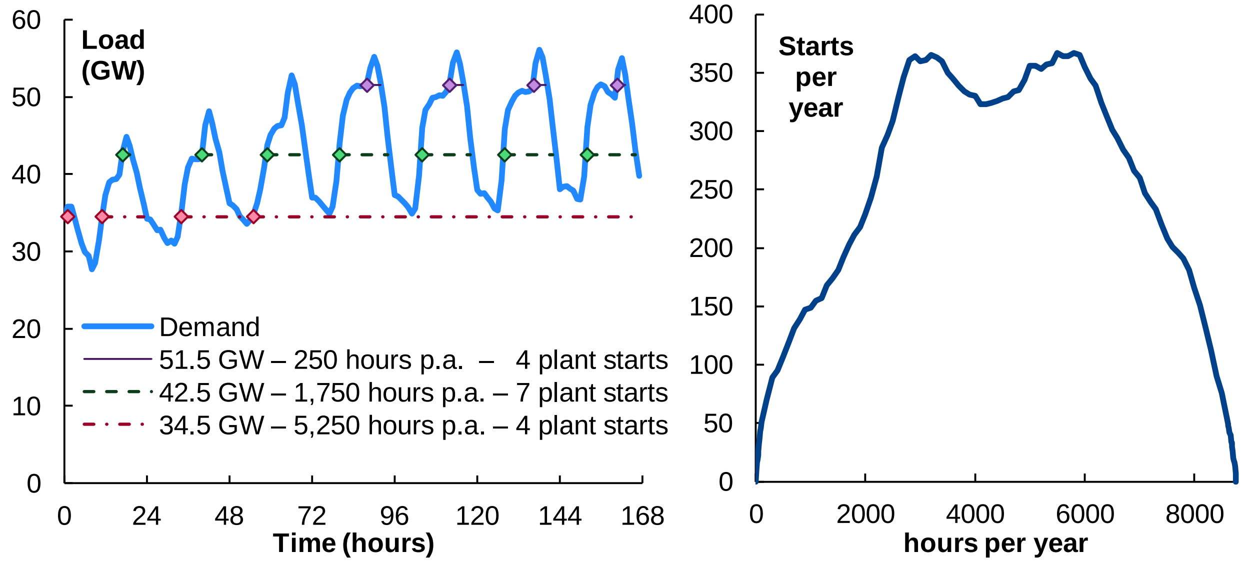

Figure 3: Demonstrating a simple algorithm for calculating the number of start-ups for different levels of plant.

This simple algorithm neglects the fact that it is often cheaper to reduce the power output of several plant than to shut down a single plant, and so will most likely over-estimate the number of start-ups. A more sophisticated algorithm that accounts for generator's decision making processes could be used in its place to gain better results.

The profile of start-ups against cumulative capacity will change over time as the level of demand changes, or if different (larger or smaller) countries are considered. We find that mapping the number of starts against the number of running hours gives a much more consistent profile over time and between countries. In order to find this relationship, the load duration curve can be used to map the number of GW of plant to the number of hours that plant is required for, just as when constructing a screening curve.

The plot to the right of Figure 3 shows this transformation with the 2011 GB data. The plant that is only required for one hour of the year (the 56.1st GW) need only start once to cover the very highest peak demand. Similarly, all plants that can run for 8760 hours (the first 21.4 GW) do not have to start up (except for maintenance). Moving from these extremes to the mid-merit plants, which are required for 3,000–6,000 hours, the number of starts increases to approximately one per day.

By considering the number of starts against the number of hours per year the plant must run, we enable this method to be translated easily to systems other than the UK. Figure 4 demonstrates this similarity for the GB system over the last 18 years (when demand has grown 18% then fallen 9%); and in an American market over the last 32 years (when demand grew by over 50%). While the shape of this curve appears to be relatively consistent within a country over time, it shows more marked differences between countries due to the different extents of electric heating and pumped hydro storage employed.

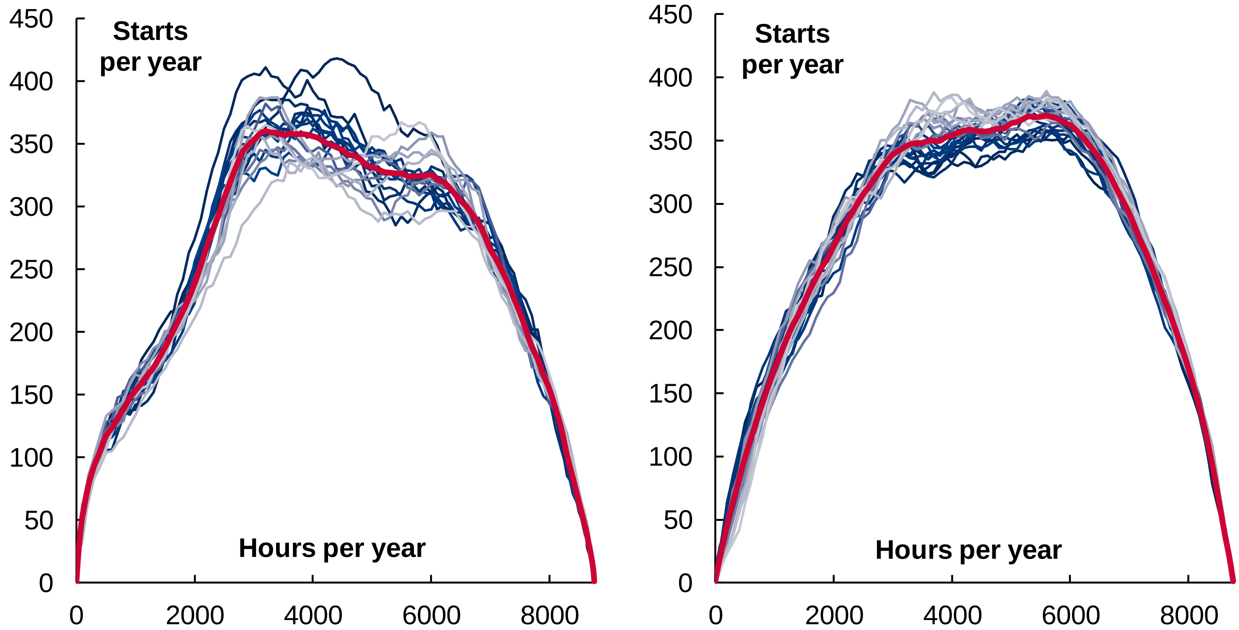

Figure 4: Average number of plant start-ups for the GB system from 1994–2011 (left), and for the New England system from 1980–2011 (right). Individual years' data is shown in blue, with the period average highlighted in red.

By combining the plant start-up costs (Table 2) with the average number of starts per year, the total annual cost of starts can be calculated for each plant type at each point within the merit stack.

Hydro scheduling

The British electricity market also includes 4.1 GW of hydro plant, which is scheduled by the GAMS model so that the energy available is sold at the most valuable times. 2.75 GW is pumped storage hydro, replenished every night; 1.36 GW is run-of-river hydro with an energy constraints which bind over a longer time period – we assume an annual constraint. In our merit order stack model, we pre-treat the demand data with a peak-lopping algorithm, as used by Borenstein and Bushnell (1999) to model California. We start by choosing a (different) level of demand for each day at which pumped storage hydro stations start to run, reducing the net demand towards this level by as much as their capacity allows. The level of demand is chosen so that the total energy used by the pumped storage plants during the day is equal to their average daily generation over the year, (9.9 GWh). After the daily optimisation is completed, we choose a similar demand level at which the run-of-river hydro starts to run, peak-lopping so that the available 4,945 GWh of energy is consumed over the year as a whole. Figure 5 demonstrates this algorithm in action: the first day has a short-lived peak and there are hours in which the full power capacity (measured in GW) of the pumped storage plants is used; the second day has a flatter peak and the full energy capacity (GWh) can be used without using the full power capacity (GW). Run-of-river hydro begins to operate at the same level during both days, as the energy constraint is not applied to each day separately.

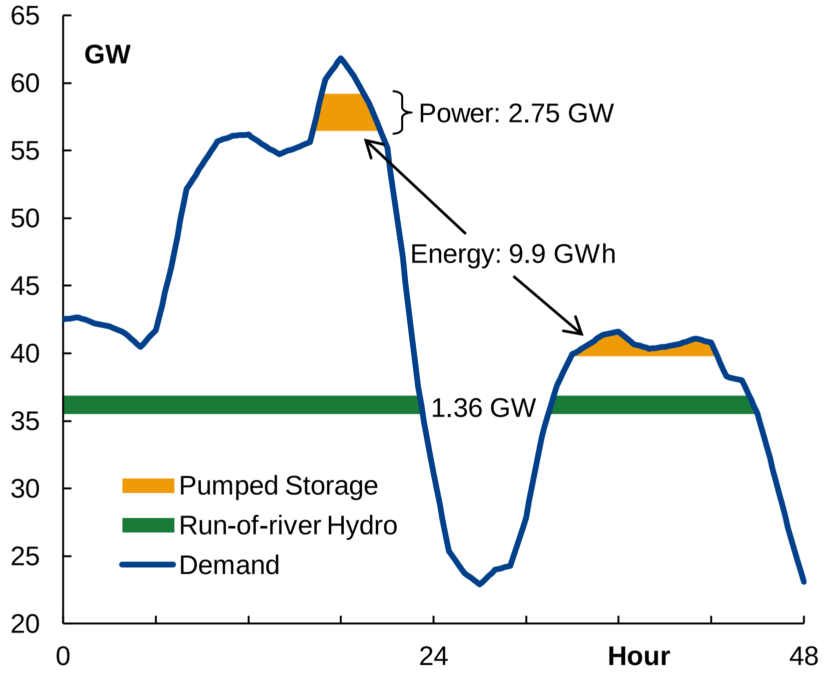

Figure 5: Demonstration of the heuristic for allocating hydro resource.

Results

Three variants of the stack and full dispatch model were chosen to highlight the most influential of the dynamic constraints. The first of these models had no constraints, and was effectively the straightforward merit order stack.

The second pair of models imposed the minimum load constraint on nuclear stations, which were unable to operate at less than 75% of their full capacity. To avoid breaching this constraint (when demand was low or wind output was high) wind power was spilled and the market price was set to minus £50/MWh, implicitly the foregone subsidy which wind stations would require as compensation.

The third pair of models both included start-up costs, while the GAMS dispatch model also included the cost of running plants at no-load, with a corresponding reduction in the cost per MWh of output actually produced (as the marginal efficiency is higher than the average).

Capacity mix

As seen in Figure 6, the long-run equilibrium capacity mix calculated by each variant of the two models was remarkably similar. The greatest deviation with the simple merit order stack was that it overestimated the optimal CCGT capacity by 2.5 GW, relative to the fully optimised dispatch result.

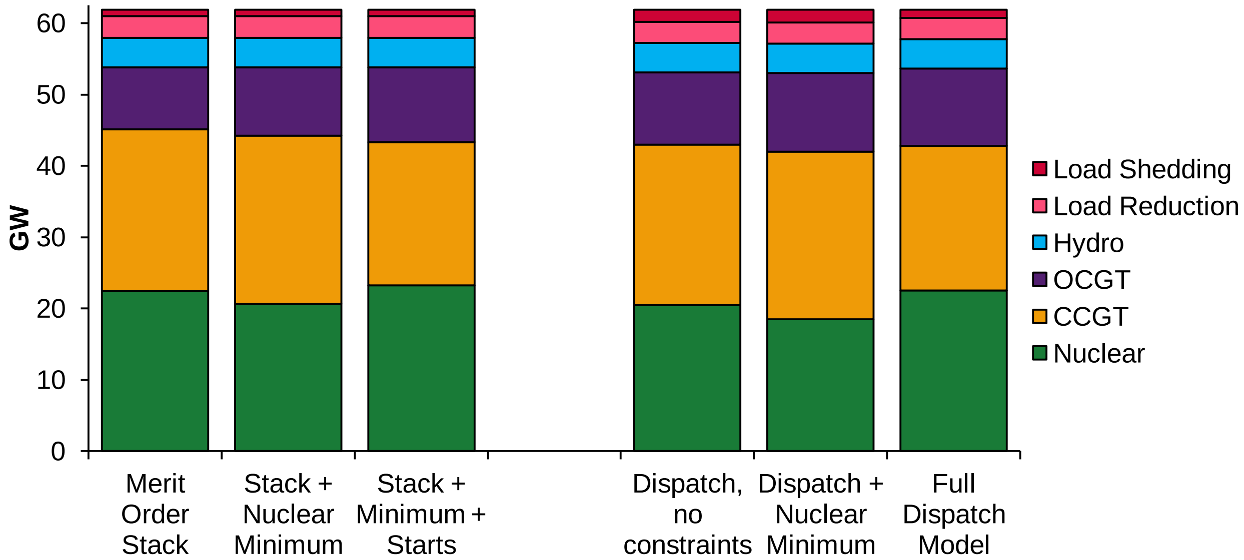

Figure 6: Optimal capacity mixes calculated by each model.

Adding a minimum stable generation for nuclear plant had the expected effect: the added cost of spilling wind meant that marginal nuclear plant could no longer compete against CCGT, so the capacity of nuclear plant fell by 2 GW with both the stack and dispatch models. The reduction in nuclear capacity was met almost equally by increases in CCGT and OCGT.

Adding the cost of start-ups reversed these changes: CCGT capacity was squeezed at both ends of the stack as it had the higher ratio of start-up cost to incremental cost[2], and so nuclear capacity increased by 2 GW and OCGT by 1 GW.

The impact of adding these constraints to the dispatch model was very similar as to the stack model, and so the net result of adding minimum stable generations and start-up costs was the same as for the stack model. The full dispatch model's result featured 1 fewer nuclear plant (500 MW) and 300 MW more OCGT capacity than the stack model with constraints. This missing capacity was made up for by increased load shedding, which peaked at 1.1 GW instead of 0.9 GW.

[2] In the case of nuclear plant, the cost of curtailing wind output was lower relative to operating costs (as this avoided the prohibitively expensive restarting of a reactor).

Outputs

The constraints and modelling methods had more noticeable impacts on the levels of plant output, as plant run times were affected by their position within the stack (in addition to the changes in capacity). Table 3 shows the energy outputs predicted by three of the model runs. It shows that the corrections for nuclear minimum output and start-up costs improve the accuracy of the basic stack model approximately two-fold.

	Full dispatch	Stack + Minimum + Starts	Simple Stack
Nuclear	192.7 TWh	198.6 (+3%)	177.5 (–8%)
CCGT	86.4 TWh	81.4 (–6%)	102.6 (19%)
OCGT	5.0 TWh	4.4 (–11%)	3.6 (–28%)
Wind spilling	–1.8 TWh	–2.0 (+15%)	–1.2 (–30%)
Load shedding	–80 GWh	79 (–1%)	79 (–1%)

Table 3: Comparison of plant outputs for different modelling methods, highlighting the deviation from the fully optimised results (in parentheses).

The greatest deviations from the fully optimised result were at the extremities of the stack: the level of wind spilling at the bottom and the levels of OCGT output at the top. In all cases, the amount of load shedding was very similar by design, as it was determined primarily by the total amount of physical capacity installed.

Prices

The impact of operating constraints is also prevalent in the prices produced by these models. Figure 7 plots the price duration curve from two of the stack models and the full dispatch model, revealing key differences in their treatment of plant start-ups.

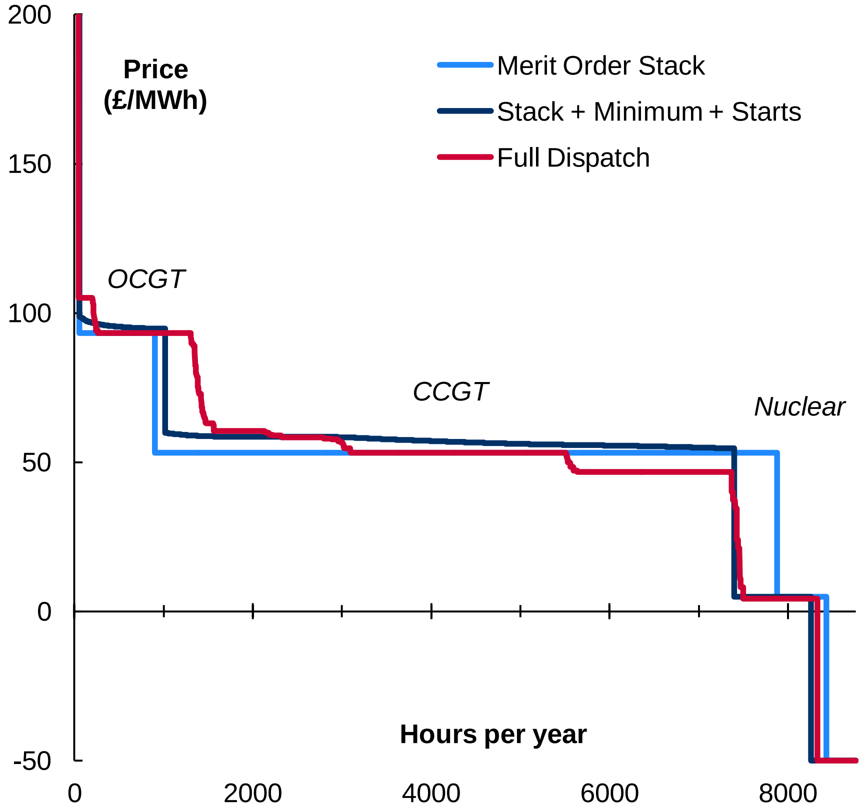

Figure 7: Price-duration curves for three of the models.

The vertical (or near vertical) segments of the price duration curve indicate transitions between different types of plant being marginal, OCGT and CCGT are marginal at approximately £100 and £50/MWh respectively, for example. The full dispatch model suggests that OCGT are marginal for 30% longer than the stack models, due to the slightly higher capacity installed and longer run-times. The transition points from CCGT to nuclear and to wind shedding are correctly predicted by the stack model when minimum outputs and start-up costs are accounted for.

In the simple merit order stack, price is determined solely by the incremental cost of the marginal plant, ignoring start-up costs. Its price duration curve therefore consists of long horizontal segments, when OCGT, CCGT and nuclear are marginal.

The modification for adding start-up costs to the stack works by allocating the cost of a start-up equally across all the hours that plant is generating. The horizontal sections of that model's curve therefore have a slight downwards slope to them, as there become more hours to spread the cost of these starts over.

The dispatch model shows a yet more complicated structure, as start-up costs are allocated only to the hours that necessitated the start-up. From 1,500 to 3,000 hours the marginal cost of CCGT is therefore higher than the incremental cost (£53.25/MWh). Conversely, in hours of lower demand during a trough (from 5,500 hours onwards) the cost of avoiding a start-up is seen, depressing the marginal price of CCGT below its incremental cost.

The behaviour of the three models can be seen more clearly in Figure 8.

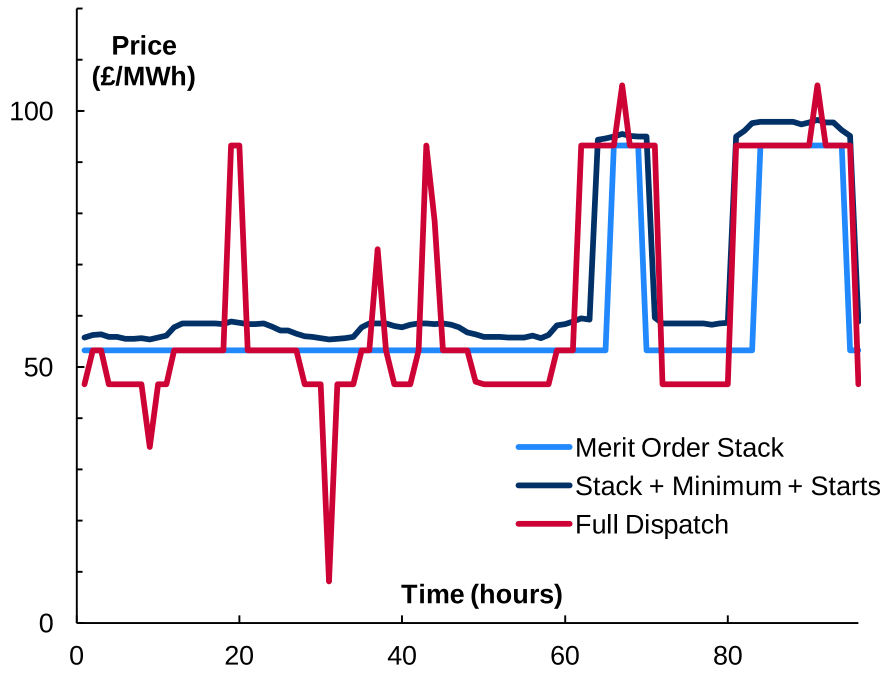

Figure 8: A sample of the price over time calculated by three of the models.

The behaviour of the dispatch model could be replicated more closely with a more complex heuristic. Two options that we intend to explore are a way to allocate the plant start-up to specific peak hours, and crediting the avoidance of a start-up during trough hours.

Conclusions

We test a merit order stack against a fully optimised dispatch model, looking at the long-run equilibrium capacity mix and electricity prices. The importance of two major constraints is highlighted: the minimum output of nuclear reactors, and the cost penalty of starting up plants.

We look at a simple method of factoring start-up costs into the merit order stack, requiring only is information on plant costs and the load profile. By inflating the cost of mid-merit plants, start-up costs move the crossover points at which OCGT or nuclear plant become the lowest cost generator, increasing the shares of these plant in the long run equilibrium mix. These changes are replicated in the price duration curve, enabling the modified stack model to approach the results of the fully optimised dispatch model.

We aim to improve the accuracy of the stack model further by allocating the start-up cost of a plant to specific hours of peak demand (the ones which made the start-up necessary), and by considering avoided start-ups in the night-time troughs.

Even when considering scenarios with a challenging level of intermittent or variable renewable generation, it seems that there is still merit in the merit order stack.

References

Arup (2011) “Review of the generation costs and deployment potential of renewable electricity technologies in the UK” London, Department of Energy and Climate Change

Battle, C. and P. Rodilla (2012) “An Enhanced Screening Curves Method for Considering Thermal Cycling Operation Costs in Generation Expansion Planning” *Universidad Pontificia Comillas*

Borenstein, S. (2005) “The Long-Run Efficiency of Real-Time Electricity Pricing” *Energy Journal*, vol. 26, no. 3, pp. 93-116

Borenstein, S. and J. Bushnell (1999) “An Empirical Analysis of the Potential for Market Power in California's Electricity Market.” *Journal of Industrial Economics*, Vol 47, No. 3, pp 285-323

Borenstein, S., J. Bushnell and F. Wolak (2002) “Measuring Market Inefficiencies in California's Restructured Wholesale Electricity Market” *American Economic Review* vol. 92, no. 5 (December), pp 1376-1405

Bushnell, J. (2010) “Building Blocks: Investment in Renewable and Nonrenewable Technologies” in (eds) B. Moselle, J. Padilla and R. Schmalensee *Harnessing Renewable Energy* RFF Press and Earthscan

Bushnell, J.B. and Wolfram, C. (2005), “Ownership Change, Incentives and Plant Efficiency: The Divestiture of U.S. Electric Generation Plants” University of California Energy Institute. http://www.ucei.berkeley.edu/PDF/csemwp140.pdf

Energy Information Administration (2010) “Updated Capital Cost Estimates for Electricity Generation Plants”. Washington, U.S. Department of Energy.

Fawcett, T., K. Lane, B. Boardman, et al. (2000) “Lower Carbon Futures for European Households” *Environmental Change Institute, University of Oxford* http://tinyurl.com/8cyyg3x

Green, R.J, and D.M. Newbery (1992) “Competition in the British Electricity Spot Market” *Journal of Political Economy* vol. 100 no 5, pp 929-53

Green, R. and Staffell, I. (2012) “How Large Should A Portfolio Of Wind Farms Be?” To be presented at the *12th IAEE European Energy Conference, Venice, Italy*.

Green, R., Staffell, I. and Vasilakos, N. (2011) “Divide and conquer? Assessing k-means clustering of demand data in simulations of the British electricity system”. Working paper available from http://tinyurl.com/c934prp

Green, R.J. and N. Vasilakos (2010) “Market Behaviour with Large Amounts of Intermittent Generation” *Energy Policy*, vol. 38, no. 7, pp. 3211-32

HM Treasury (2011) “Carbon price floor consultation: the Government response” London: HM Treasury http://tinyurl.com/96dwppr

International Energy Agency (2010) “Projected Costs of Generating Electricity”. Paris, IEA, NEA, OECD.

Joskow, P.L. and E. Kahn (2002) “A Quantitative Analysis of Pricing Behavior in California’s Wholesale Electricity Market during 2000”, *The Energy Journal*, vol. 23, no. 4, pp. 1-35

Kirschen, D and G Strbac (2004) “Fundamentals of Power System Economics”, Chichester, Wiley

Lamont, A.D. (2008) “Assessing the long-term system value of intermittent electric generation technologies” *Energy Economics*, vol. 30, no. 3, pp. 1208-1231

Mills, A. and R. Wiser (2012) “Changes in the Economic Value of Variable Generation at High Penetration Levels: A Pilot Case Study of California” Report LBNL-5445E, Ernest Orlando Lawrence Berkeley National Laboratory

Monopolies and Mergers Commission, (1981) “Central Electricity Generating Board: A Report on the Operation by the Board of its System for the Generation and Supply of Electricity in Bulk”, HC 315, HMSO, London

Mott MacDonald (2010) “UK Electricity Generation Costs Update”. London, Department of Energy and Climate Change.

National Grid, 2011, “National Electricity Transmission System Seven Year Statement”. Coventry, National Grid plc.

Parsons Brinckerhoff (2011) “Electricity Generation Cost Model - 2011 Update Revision 1”. London, Department of Energy and Climate Change.

Rosnes, O. 2008, “The Impact of Climate Policies on the Operation of a Thermal Power Plant” *The Energy Journal*, vol. 29, no. 2, pp. 1–22.

Stoft, S.E. (2002) “Power System Economics: Designing Markets for Electricity” Chichester, Wiley

Strbac, G., Aunedi, M., Pudjianto, D., *et al.* (2012), “Strategic Assessment of the Role and Value of Energy Storage Systems in the UK Low Carbon Energy Future” *Carbon Trust.*

Usaola, J., J. Rivier, G. Sáenz de Miera, M.Á. Moreno, and M. Bueno (2009) “Effect of Wind Energy on Capacity Payment: The Case of Spain” 10th IAEE European Conference, Vienna