The characteristics of electricity storage, renewables and markets

Michael Waterson[1]

University of Warwick and ZEW, Mannheim

Revised October 2016

Abstract

This paper accepts the widespread view that as electricity generation systems transition towards a greater proportion of renewables provision, there will be an increasing need for storage facilities. However, it differs from most such studies in contrasting the private incentives of a storage operator with the public desirability of bulk storage. A key factor in the context of a market such as Britain, where renewable energy largely means wind generation, is the nature of wind generation itself. The problem of wind's high variance and intermittent nature is explored. It is argued that not only is there a missing money and a missing market issue in providing secure energy supplies, there is also a missing informational issue. A key opportunity for new storage is participation in a capacity market, if the setting is right.

Keywords: Electricity storage; Inter-temporal arbitrage; Renewable generation; Capacity auctions.

[1] This paper reports on work carried out in part with John Barton, Lisa Flatley, Monica Giulietti, Luigi Grossi, Robert MacKay, Murray Thompson and Elisa Trujillo. I am particularly grateful for Elisa's comments on a previous draft, as well as the comments of two anonymous referees. Much of the work reported on here has been funded by the EPSRC under grant EP/K002228 to the IMAGES project. Author's email: Michael.waterson@warwick.ac.uk

1. Introduction

This paper examines aspects of energy storage from the viewpoint both of opportunities for private firms and what would be socially desirable. The context in which this is set is important. Claims about the generation cost per MWh of renewable-generated electricity being competitive with more conventional power plants are commonly made. In Britain, there is a plausible argument that the levelised cost per MWh of onshore wind is comparable with the price the Government has agreed to pay for a new nuclear plant under development (DECC, 2013, chart 3). Of course these costs are not equivalent in terms of supplying power as required, since both have limitations. Solar power is intermittent; wind power is both intermittent and has an extreme and time-correlated variance. Nuclear power, on the other hand, is inflexible. Thus both renewables and nuclear power require additional facilities to be present in the electricity delivery system to facilitate correspondence between demand and supply. Traditionally, these have taken the form of more flexible fossil-fuelled generating plant on grid. However, if the presence of renewables increases, and at the same time, the role of fossil-fuelled generators decreases due to retirement of ageing or overly polluting plant, additional measures will be required. One obvious additional facility is increased energy storage (Denholm et al, 2010; DG ENER, undated; Greve and Pollitt, 2016).

The main novelty in this paper lies in the comparative analysis of market-based and socially desirable storage (largely ignoring power quality operations such as frequency regulation that take place very near to real time). The main finding in relation to market arbitrage-based storage is that diurnal storage is currently the obvious source of profit in Britain, given the large diurnal price differences, the relatively small price differences over days, and the unpredictability of wind over more than a short period. Such arbitrage activity is most suited to storage over no more than a few hours, implying the leading technologies are likely to be heat storage or compressed air technologies (on which see later). However, there are clear social benefits to longer term storage based upon saving peak generation and a reduced need to curtail renewables, which point more towards compressed air energy storage. But these benefits, we argue, would not be captured by the store under arbitrage mechanisms, because they require look-ahead times much longer than available given current weather forecasting and market pricing models that currently do not exist. There is extremely limited information, and missing markets.

This modifies the emphasis on market issues compared with Newbery (2016), who points to the problem of missing markets, but also to missing money. We add to these the issue of missing information, which renders the development of some markets extremely difficult. The essential difference between missing markets and missing information is the following. Missing markets are those that could exist, or would exist, under an alternative framework, for example given a different regulatory structure. Missing information is something that prevents a market existing, because there is scant or no information on which to form expectations of the future and hence to formulate prices. For example, if there is no way of knowing whether it will be extremely windy or extremely calm next week, expectations cannot be formed.

We further argue that interconnectors and smart metering, alternative mechanisms for bringing demand and supply into line, are for different reasons not obviously suited to these longer-term issues. Moreover, it is unlikely that prices are able to provide enough of a signal to storage operators looking to longer-term storage. Therefore, the obvious alternative to attract longer-term storage appears to be a capacity mechanism. But in turn this relies on the treatment of storage relative to other forms of capacity, given that storage is currently viewed both as a consumer and a generator and hence pays charges related to both, and also is limited in that it cannot commit to an indefinite supply time. The paper sets out some of these arguments. Our context is Great Britain, but there are some parallels in other countries. However, the requirements for storage will differ between countries somewhat, dependent on the make-up of renewable generation and the pattern of demand across those countries.

The paper proceeds with a brief resume of storage technologies and discusses the characteristics of renewables, proceeds in section 3 to consider the role of electricity pricing and arbitrage, then subsequently discusses the differences between private and social benefits to storage in the context of uncertainty about the amount of renewable power in section 4, goes on to consider alternative market forms in the light of this in section 5, and finally discusses the policy implications of the analysis in section 6.

2. Storage technologies and characteristics of renewables

There is a widespread view (e.g. Denholm et al, 2010; Evans et al, 2012) that as countries increase the role of renewable resources, principally wind and solar, in electricity generation, storage will become more important in the role of balancing supply and demand for electricity. The operative issues are how much of which type will be required, what benefits could it bring, and how will it make money. What storage is "required" and how it makes money are separate issues[2]. This links with the proposed technology employed to "store" electricity and the institutional framework into which it slots.

The electricity industry in England and Wales (and to a lesser extent, Scotland) is extremely vertically disintegrated so far as the role of the grid operator is concerned. Specifically, the grid operator cannot own storage facilities, a situation which contrasts with that in Italy, for example, where Terna is active in developing storage. Moreover, there is no equivalent of the regulated electricity utility, as exists in large parts of the USA. Therefore, it is not possible for the grid operator (nor to a large extent the distribution companies) directly to internalise the benefits of storage that accrue to themselves. Nor has investment in storage been mandated, as it has for example in California. In

[2] Much of the literature's focus (see e.g. Cavallo, 2007) has been on system requirements, without regard to how storage will be provided commercially. From the opposite standpoint, Zafirakis et al (2016) examine the potential commercial role of arbitrage, without regard to its social desirability.

Britain, external market mechanisms through contracts with National Grid, if present, are the method by storage may be incentivised. More generally, many systems including the British electricity system, suffer from storage not being accommodated within the standard framework of generator or supplier; indeed it is commonly treated as both, rather than as a facilitator for smoothing capacity issues. The role of storage in particular of batteries in primary and secondary (within seconds) frequency control is recognised, as also for example in Germany (Energy Storage Update, 2015), but very little attention has been paid to longer-term storage in other technologies such as CAES.

Batteries are good at fast reaction but short duration requirements, and so are best suited to quality maintenance such as frequency regulation- for example in Britain, plans are under development to tender for facilities capable of reaching full power in less than one second to be sustained for some minutes and this is attracting interest from battery storage. Anything over a few minutes is not suited to the use of batteries, which at least currently have very limited energy capacity. Hydroelectric storage has the capability to supply electricity at significant rates over a relatively long period, and it can reach maximum power in a matter of seconds, but it is limited in the extent to which it can be constructed economically, because it relies on the existing geography of the country.

Heat-based storage (e g HTTP; High Temperature Thermal Power) and compressed air electricity storage (CAES) have the additional potential to store electricity to be delivered for a period of a few hours to over several days. It is these latter two that are most related to the subject of this paper, given the diurnal pattern of prices. Heat-based storage has relatively low fixed cost but high operating cost compared with CAES and so is best employed for relatively shorter periods. However, it should be emphasised that fixed costs and running costs for each of these technologies, particularly the newer ones, are subject to significant uncertainty and a degree of optimism on the part of their proponents, together with a certain amount of commercial secrecy. Therefore it does not seem reasonable to give more than these indications regarding the relative cost characteristics of the various storage technologies.[3]

Implicitly, the discussion below is framed with CAES-based storage in mind, although it can be applied to other technologies. Therefore it may be useful to explain the technology briefly. CAES involves compressing air into a store, for example a cavern, using electrically-driven pumps. When power is needed, the operation goes into reverse; air released from the cavern drives turbines which produce electricity.[4]

The key intermittent renewable resources being employed to generate electricity at grid scale in Britain are wind power, solar power and prospectively tidal power.[5] None of these is biddable in the way that conventional combustion power plant is. Tidal power is intermittent but almost completely predictable. Solar power is intermittent but relatively predictable. In particular, we can predict that it will not be available at night! Wind is intermittent but arguably relatively unpredictable except in the short run, and erratic. Wind, significant both onshore and offshore, is by far the largest proportion of generation in the UK under current circumstances, and is likely to remain so. Hence this is the prime focus of the analysis.

For the purposes of analysis, we use data for Great Britain from the period end-November 2014 to end- September 2015. The start date is determined by availability of data for wind forecasts, which have traditionally only been displayed temporarily on the BM Reports website provided by Elexon. Since end-November 2014, the Gridwatch site has been recording the data feed including wind

---

[3] More detailed information regarding comparative capital costs are provided in Evans et al (2012), but recent attempts to construct commercial HTTP plant in Britain have run into problems suggestive of over-optimistic assumptions on costs.

[4] There are two variants: adabiatic CAES requires no additional fuel, whilst existing facilities such as at Huntorf in Germany, involve using the air in a gas turbine; of course air when compressed warms and when decompressed cools, and in both cases the technical design needs to accommodate this.

[5] Other renewable sources include biomass and run-of-river plant (of which there is very little in Britain), but in any case these are more controllable and therefore do not lead to the problems discussed in this paper. There are also relatively limited storage facilities based on hydro-electric plant, because the terrain is largely unsuitable.

forecasts.[6] In illustrating features of these data, we sometimes choose shorter periods for clarity of presentation, but the main analysis is done using the 10 month period.

Table 1 characterises the wind generation pattern. Whilst on average over 2.5GW are being generated currently, the variance is very large, so that on a high proportion of occasions less than 1GW is being generated. In Britain, and many other countries, one characteristic is that the weather can be calm for several days in a row, so that wind generation suffers significantly from needing some form of backup. Figure 1 illustrates a recent three month winter period drawn from the same underlying data, which exhibits this point very clearly; there are several periods in which generation is below 1GW, although the mean generation is approximately 3.5GW in those winter months. Moreover, there is some bunching of low generation levels in the series at particular points.

In order to get a handle on the bunching characteristic of the data, we analysed periods of a day or more in which generation was less than 1GW for each observation. There were 17 such incidents in our 10 month period, the longest of which lasted for five days. Average generation across each of these incidents is less than 600MW. In other words, if wind were to be relied upon to generate at least 1GW, ancillary facilities (for example storage or back-up generation) would need to provide almost 10GWh over a 24 hour period in order to maintain this requirement.

Of course, significant wind plant is under construction, particularly in off-shore areas. However, the existing wind plant is spread around the country already and there is no particular reason to believe that the time properties of wind generation will experience a reduced variance once a greater capacity is installed. If anything, the move to more off-shore production will increase uncertainty. Currently Ireland, which has around double the wind generation capacity installed in proportional terms, experiences very similar patterns of wind generation.

[6] Its recording of other variables has a longer history. Intending users of this site should note that the series do contain several unmarked gaps, mostly of a few hours' duration, meaning care needs to be taken in creating a consistent continuous time series.

In southern Europe and North Africa, where the weather consists mainly of sunny days, storage can take the form predominantly of storage for a few hours to move electricity from the middle hours of the day more towards the evening. Spain has several examples of heat storage plants, for example involving molten salt, which fulfil this task. However, in the case of wind generation in Britain, storage that is capable of moving generated power within the day will not be sufficient to tackle the sustained calm periods.

What about wind forecasts? Figure 2, which for clarity takes a two-week snapshot from our data, compares the initial forecast, produced around 36 hours prior, to the final wind out-turn. This is sufficient to illustrate the typical pattern, which is that short-term forecasting of wind generation is good on average. Knowing reasonably well how much electricity wind will generate tomorrow is not a major problem, although capturing the precise timing is more problematic. Over our complete sample, the unweighted average proportionate absolute error per hourly observation is 0.38 for the initial forecast, but this improves to 0.26 for the final forecast. Moreover, although precise timing is slightly awry, the closeness with which forecasts capture the overall pattern is apparent[7]. Having said this, there are clear issues with current wind generation forecasts. One is that "mean" forecasts are actually over-optimistic. A second is that the mean is not the only, or even perhaps not the most relevant, statistic of relevance to system stability. Third, capturing wind generation patterns over a longer period is markedly more difficult, as our investigation of the Belgian system later in the paper reveals.

3. The role of pricing and arbitrage

As is well known, wholesale spot electricity prices vary significantly over time, whereas supply to final consumers is normally at prices fixed in advance. Therefore since they face significant price risk, suppliers engage in a portfolio-building strategy of signing contracts to buy at various points in the

[7] Some other scholars, for example Forbes and Zampelli (2016) are much less positive about short-term forecasting of wind generation than we are here. In fact, their view if anything strengthens the relevance of a point we make later in the paper, concerning missing information.

future (for example, year-ahead, week-ahead, day-ahead), only purchasing a very small amounts at spot prices. Because demand fluctuates over the 24 hours of the day, it is normal in developed countries for spot and near-time prices to be significantly higher in peak periods, a few hours in the morning and some hours in the evening, plus perhaps (dependent upon habits) some peak hours in the middle of the day in addition. These prices are low overnight, though. This means that there is significant commercial potential for storage of wind to smooth daily peaks. Figure 3 illustrates the pattern of half-hourly spot market prices over a period of around six weeks from within our data, in which the vertical gridlines are purposely at 48 half-hour intervals. The diurnal pattern to prices is clearly visible, and this is the most obvious tendency in the data. Moreover, there is a half-hourly day-ahead market for electricity which facilitates inter-hours trade. On the other hand, figure 4 illustrates the relatively narrow gap between seasons of the year in mean electricity prices, something which contrasts strongly with the pattern in gas prices, which do not vary much within the day, but vary significantly across seasons.[8]

All this suggests that arbitrage within the day has the potential to be a profitable activity for a store, provided that it is sufficiently efficient. Arbitrage across longer periods of time is much less likely to be profitable. Figure 5 drills down more deeply into the diurnal variation, using the example of the week beginning Monday 12th January, 2015.[9] The point of the figure is to note that the diurnal pattern of prices incorporates typically a few hours in the night where prices are rather low, and a few hours in the day when they are rather high. This is important because the average gap between the peak and the off-peak price in the market is not sufficient to make a simple strategy of buying off-peak and selling peak (forward) profitable, most of the time. In fact, we have confirmed this to be true for recent years' data relating to peak versus off-peak prices[10].

[8] Of course gas does not have the property that supply must always equal demand, because pressure in the supply pipes may be varied, and gas is used significantly more in winter for heating purposes.
[9] There is nothing special about this week; it has been chosen randomly.
[10] Unpublished research by Giulietti, Grossi and Waterson.

However, our analysis (using current prices rather than projected future prices) suggests that there is scope for profitable arbitrage operations which take advantage of a few hours of reliably lower prices at night and a relatively short period of higher prices during the peak hours of the evening to operate a store on a diurnal basis; the basis could for example be early morning purchase at spot, backed with resale forward to later in the day making a clear temporal arbitrage. Obviously, the extent to which this is profitable depends upon several factors, including the efficiency of the store, the number of hours over which it charges and discharges, and the size of the store (through its lesser or greater effect upon clearing prices). The calculations are most straightforward if we assume that the store's operations neither affect market prices nor the strategies of other parties in the market such as generators (see Flatley et al, 2016a, but also Hutchinson, 2015).

Our calculations (Flatley et al, 2016a) show that a modest store in excess of 70% efficiency with capacity equal to up to five hours' worth of input pumping is likely to make an operating surplus, assuming that price differentials remain more or less consistent over time.[11] More extensive facilities would lead to a decline in the differential and therefore reduce profitability, as well as the marginal utility, of such facilities.[12] [13]

These calculations, other work we have pursued (Flatley et al, 2016b), and indeed figure 5, show there may be some potential for profitable longer-term storage, extending beyond one day. In the figure's example, it is clear that Saturday is not a profitable proposition for storage, but that storing in the early hours of Friday to discharge on Sunday evening would (in this particular week) be worthwhile. However, to do this the store would need to be proportionately larger and therefore

[11] Some capital subsidy might still be required, dependent on likely construction costs.

[12] Whilst it is commonly assumed that diurnal price differentials will increase over time with increased renewable deployment, there are several potential reasons why this may not happen, and indeed in Germany we observe a narrowing of differentials in recent years (see Grossi et al, 2014).

[13] The nature of the strategy pursued will also matter. Cruise et al (2015) assume a moderate degree of sophistication in the strategy towards storage. Both more and less naïve strategies are possible. A simple strategy of storing and delivering every day would not, according to our calculations, make sense, particularly in summer months.

more expensive, and the store's efficiency becomes of crucial importance in influencing the possibilities.[14] It would also need to have accurate forecasts of prices over a period of a few days. Figures 6a and 6b, drawn from Flatley et al. (2016b), illustrate the point. With a very efficient small store, beyond current capabilities but at envisaged new generation efficiency levels (Evans et al, 2012), privately-optimal operation remains very short term, usually diurnal, but a similarly efficient store with a much higher capacity to power ratio would optimally retain energy in storage for longer.

It is important to note that these points are made in the context of not knowing the likely cost of construction of such a store; they assume that it is available. Further analysis is being undertaken to gain a better perspective on this. But clearly, if something is unprofitable because variable costs exceed present revenues, then even without taking construction costs into account, the business model will not stand up as regards profitable operation.

4. Consideration of private versus social benefits in the context of uncertainty

An obvious question is whether entrepreneurs, seeking private profitability by investing in storage, will achieve social optimality. To put this another way, do spot (and near-dated forward) prices provide a sufficient signal to the market? This is a priori unlikely, since there are external benefits of storage that devolve to other parties, such as network operators, which the store cannot easily capture. The social benefits include in principle: meeting peak demand, thereby saving capital expenditure on new peaking plant (at the expense of increased storage construction costs, of course); avoiding some curtailment of renewable energy; reduced expenditure on grid reinforcement; fuel saved through reduced ramp rates; reduced need for low efficiency plant to operate, etc. (only these last two of which are likely to be represented in wholesale prices, and

[14] Therefore, since storage technologies differ in efficiency, the technology adopted will have an influence on the pattern of operation, although similar simulations with a 90% efficient plant show a rather similar overall pattern to those in figure 6a.

therefore capturable through arbitrage). Essentially, there is a missing market problem due to uncaptured positive externalities.

Work by Barton and Thomson (2015) has calculated that under certain assumptions there is a cost-benefit analysis case for storage up to around one week to 10 days ahead. The precise length of time depends on the particular scenario adopted, but 14 days is an outside estimate. Therefore, inter-seasonal storage is clearly ruled out on cost-benefit grounds given current or foreseeable technology, but storage substantially in excess of one day would be socially worthwhile, on the basis of these calculations. Figure 7 is drawn from their work. This analysis is based on social operation rather than private profitability considerations, and private companies would probably wish to stop well short of this, in the absence of subsidies. However, this work does incorporate one specific assumption that is extremely favourable to storage, in that it assumes future wind power fluctuations are perfectly observable; there is no uncertainty of any kind in the modelling. Hence storage is assumed to be optimally deployed. Nevertheless, it provides some sort of benchmark.

Why does uncertainty matter? Wind generation output is inherently variable. Also, once we look more than a short period ahead, predictions become subject to very large errors. Analysis performed on Belgian data[15] (where wind generation predictions are released daily for the week ahead) shows the extremely wide range of possibilities once we get to the outer range of this prediction- for example on a particular day a week ahead, wind might generate at 5 or 50% of available capacity, clearly too wide a range to be useful for planning a store's strategy. Figure 8 shows a small snapshot in terms of percentage of capacity predicted to be produced as against percentage actually produced, where the predictions are those produced from one to seven days ahead.[16] As in Britain,

[15] Belgium is a small country, with a small stock of wind generation. Nevertheless, its predictions provide a useful pointer, because Belgium experiences similar weather to Britain, and has both onshore and offshore facilities. The period studied is short, but provides sufficient evidence to make the point.
[16] The network operator website removes predictions a day at a time. Hence, tracking them over time requires capturing predictions when they are first made, then matching this to the day for which they are made.

forecasts produced one or two days ahead are able to forecast generation moderately accurately. However, as we move towards predictions for longer ahead, it can be seen from the figure that two things happen. First, the predictions become worse, but second, they do not necessarily move nearer to the outcome as time gets closer. Thus operating a store to absorb these fluctuations becomes inherently difficult. Suppose the store is half full and it is currently not very windy. Is it worth starting to discharge the store now? Well, this depends in general on how windy it will be in the future. If it is getting windier, then it may be worthwhile. However, if it is going to be dead calm tomorrow, or the next day, then it is probably not worthwhile discharging at present, but rather waiting.

So, developing and operating a store that has the capacity to hold several days' supplies is socially beneficial but not commercially sensible in terms of arbitrage in the absence of a sophisticated strategy, and possibly not even then. (Contrast this with a diurnal store that charges at dead of night and discharges the next day in the evening peak; nothing very sophisticated about that).[17]

Moreover, there are preconditions for storage to be able to operate through price signals. Buying in the night and selling peak later that same day is feasible because a future price exists. The optimal capacity of storage to build (in a country like Britain which will be increasingly dependent upon wind energy) inherently relies on when, in terms of lead time, you can make decisions about how to operate it, regardless of whether it is operated commercially or in the public interest. So a question arises as to how much information about future wind generation, as captured by forecasts, is built into prices.[18]

[17] At the other end of the spectrum, a storage system capable of smoothing output over a year would be enormous (Esteban et al, 2012) and would be socially excessive in the case of Britain.
[18] A related question is the length of look-ahead period required. We have done work on this (Flatley et al, 2016b), but it is not reported on here.

Potentially, markets are good at absorbing uncertainty about the future. However, the agents in such a market need information in order to form expectations about what the market will bring forth. In terms of wind forecasting, there is a clear problem that because the range of possible scenarios is so broad even a few days ahead, there is missing information likely leading to a missing market.

A preliminary question here is whether wind forecasts currently influence market prices, for if they do not, then it is unlikely that longer-term factors will be captured in price information. Consider predicting price based upon load and wind generation, in order to examine whether there is a pattern. The implicit assumption in this analysis is that, apart from wind, other generation is essentially controllable and will involve a choice amongst gas and coal generation according to factor price movements. Wind, on the other hand is not controllable, but may be forecast. Of course, load is potentially endogenous, because load and price may be viewed as being jointly determined by demand and supply factors. Hence we instrument for load in examining the impact of wind forecasts and load on price. More formally, our model is as follows:

$$p^f_{t+1,t} = g(L^f_{t+1,t}, W^f_{t+1,t})$$

where the meaning of $x^f_{t+1,t}$ is the forecast at *t* of the value in *t*+1, *p* is price, *L* is load and *W* is wind generation. A similar model could be used for longer-term forecasts, though see below. The purpose of the model is to see whether what we know one day ahead about wind provides useful information about prices. Load is a good first order determinant of price in a system in which all generation is biddable, and load forecasting has a long history.   In a system with substantial wind, this will influence the amount of biddable wind required and hence price.

Load is instrumented for these purposes by a deterministic polynomial model of the day and time of day, things that are inherently knowable and match demand well but clearly not endogenous. Wind forecasts come from the Gridwatch site and price is the APX day-ahead forward mid-price for that half hour (or, alternatively, and with very similar results, the weighted average for the half hour, a two-period block and a four-period block styled RPD). The IV regression run on our ten months of hourly data is shown in table 2 below. Both variables are very significant.

Based upon this regression, we find that the elasticity of price with respect to load, at the mean, is around 0.6. The elasticity of price with respect to the wind forecast is just over 0.1. To interpret this, note that a one standard deviation change in the forecast for wind, at the mean, would result in a change in price of £1.63, where the average price is just over £40 per MWh. Interpreting this figure in the context of our question, near-term wind forecasts clearly do have an impact on price which is measureable. At the same time, the same sort of exercise looking for impacts on week-ahead prices[19] shows no correlation. Hence a significant calm period would appear not to change price by enough to make it worthwhile to store for longer than a day, at current values. This is because the diurnal variation in price as a result of demand patterns dwarfs the impact of wind on price. Of course, in future as the penetration of wind grows, its impact of price is likely to grow, but even an effect double or triple the magnitude of that captured here is modest in comparison to other factors.

This leaves open the question of possible longer-term forecasts. The 36 hour forecast reported via Gridwatch has a similar impact on price in a similar regression, somewhat smaller in magnitude although still well-determined (but with a lower R-squared). But what of longer-term forecasts that might predict calm or a windy period for several days? Here there is a clear case of missing information- these forecasts have extremely limited predictive value.

Therefore, a policy of purchasing spot and selling a (say, one week) future is not going to work if those traders operating in the futures market have literally no idea where the price is going to be next week, because they don’t know how windy it will be. Thus, forward prices beyond a day or two

[19] Week-ahead prices, or rather price assessments, were obtained for this purpose from Platts.

will not be influenced by future wind generation, given current generation technology and market developments. At present, standard futures contracts for, say, one week ahead are not differentiated by hour; they imply constant supply over the week. It is possible of course that more finely defined contracts will develop. But the lesson of the Belgian wind analysis is that, unless wind forecasts somehow improve markedly, there will be no reason for such forward prices to develop, because the market participants will have so little information regarding future wind on which to draw in forming their expectations.

5. Alternative market forms and solutions

Conceptually there is a more straightforward answer to the desire to bring private and social incentives into line by tackling the external effects of a more intermittent electricity generation system.[20] This is to operate through a capacity storage market capable of commanding several days' supply in the event that the wind is blowing either too little or too much to be able to generate an amount of electricity that can be accommodated on the system without violating its constraints. Such a market would be capacity not for an hour or two in the event of sudden tightness of supply, but rather for gradual operation over a period of a few days. Currently, this is a potentially missing market, because the time scale over which the capacity market in Britain might require generation has not been specified.

The capacity market auctions in Britain have not successfully attracted novel forms of storage within the clearing price. One plausible reason for this is the technologies' potentially uncompetitive nature. In part this may be because storage is arguably poorly treated in the current network

[20] In the much longer term, when the current contracts under which renewable suppliers are largely de-risked, there is the potential for renewables to be supplied at what is the current market price. This would increase the incentives for investment in storage directly connected with wind and solar farms.

operating codes, being treated both as a consumer, and also a generator, therefore paying grid charges related to both. The other possible reason is that, despite an attempt being made to design the capacity auction in a fuel-neutral manner, the open-ended commitment required on being called to produce is much more difficult for a store to guarantee than for example it is for a diesel generator. Alternatively, thinking more laterally, there is the possibility for a portfolio bid, consisting of a storage facility plus a higher running cost backup facility (such as a "diesel farm") making a joint bid. This would mean that the initial capacity would be guaranteed by the storage facility, with the diesel farm only operating in the unlikely event that the capacity is required for an extended period. To our knowledge, no such bids have been received, nor are they allowed within the current code.[21]

The discussion above is premised on the need for some form of storage in order to smooth fluctuations in wind. Whether this is true needs more motivation. There are at least three challenges. The first is that the times when the weather is calm are few. This challenge lacks substance in that, as we have seen, relying on even 1GW from the current set of wind plant is problematic. Moreover, it ignores the point that periods of calm tend to come together, rather than being scattered randomly across time. Hence it is much more difficult to shift load than if it merely meant moving it to later in the day.

The second challenge is that interconnectors can suffice to deal with the problems of calm days. Interconnectors certainly have a role to play, as Newbery (2016) argues forcibly. However, to the extent to which there is inter-country correlation in wind speeds, this does not bode well for the use of interconnectors, since all countries will be bidding for generation capacity from their neighbours.

[21] Implicitly, they appear to be ruled out through Rule 2.3.4, p.27 in OFGEM (2014). A more far-reaching suggestion is made and developed through some theoretical examples by Greve and Pollitt (2016), who suggest that a Vickrey-Clarke- Groves auction mechanism be adopted in which companies providing services to the system operator could bid packages of services into the system.

The third challenge is that demand management can tackle the problem. Again, this ignores the fact that calm periods tend to be bunched. It also ignores the fact that, if it is calm in the evening, unless people wish to sit in the dark, consumption will be needed. Whilst the use of washing machines and the like can be postponed for some hours, lighting demand and demand for leisure purposes (viewing television, surfing the internet) is not the same. In fact, much of what is considered currently as demand-side response is simply a supply-side response (i.e. generation) on the other side of the meter.

6. Concluding remarks and policy implications

Our subject has been the role of longer-term storage in tackling the increased uncertainty inherent in an electricity system with a higher proportion of intermittent renewable generation. Why does the topic matter more generally? Although the question of the technical extent to which storage can be employed is important, storage faces limitations of a more economic nature which will curtail the degree to which it is practical to develop storage given current technology, beyond its uses for very short term quality maintenance.[22] Importantly, it creates positive externalities that are currently not captured. This is clearly a key element in evaluating storage possibilities from a system-wide viewpoint. But it is evident that in order to progress the economic evaluation, the characteristics of renewable generation, the characteristics of storage, and the subsequent need for developments, for example in wind generation forecasting, are key. There is missing information in addition to missing money and missing markets, and the missing markets are not easily developed in the presence of the lack of information in wind forecasts beyond the very near term.

[22] Storage, in whatever form, has a potential (additional) role for example in what is called secondary firm frequency response within the National Grid system, but this is not our concern in this paper.

There are two main policy implications. One is that for storage to develop and compete in the capacity market, some reform of that market should be developed, namely to allow portfolio bids that encompass more than one technology from a single bid. This would enable storage to compete on more equal terms with other technologies capable of providing capacity into the market. Moreover, it would be a step along the way to a more holistic approach (Greve and Pollitt, 2016) to offering tenders for the various services needed to keep the system running smoothly. This might include tendering for a suite of services including short term quality maintenance and capacity assurance.

The second implication is that increased deployment of wind generation exposes the need for more accurate and longer-term forecasting of wind generation. This means a refocus towards greater granularity so far as off-shore forecasting is concerned, in addition to forward prediction of periods when wind will not be able to generate. Forecasting the mean wind speed is not as important here as focusing on characterising the tails of the distribution, a topic that I leave for further work.

References

Barton, J P and Thomson, M (2015) "High-Temporal Resolution analysis of UK power system used to determine the optimal amount and mix of energy storage technologies", Presentation to UK energy storage, November 2015.

Cavallo, A. (2007) Controllable and affordable utility-scale electricity from intermittent wind resources and compressed air energy storage (CAES), *Energy* 32, 120-127.

Cruise, J Flatley, L, Gibbens, R and Zachary, S (2015), Optimal control of storage involving market impact and with energy applications, arXiv 1406.3653v2.

Denholm, P, Ela, E, Kirby, B and Milligan, M. (2010) The Role of Energy Storage with Renewable Electricity Generation, National Renewable Energy Laboratory, Technical Report NREL/TP-6A2-47187, Colorado.

Department of Energy and Climate Change (2013) Electricity Generation Costs 2013, DECC, London.

Energy Storage Update (2015), Germany's power market reforms provide platform for storage growth, http://analysis.energystorageupdate.com/germanys-power-market-reforms-provide-platform-storage-growth Accessed October 2016.

Esteban, M, Zhang, Q and Utama, A (2012), Estimation of the energy storage requirement of a future 100% renewable energy system in Japan, *Energy Policy* 47, 22-31.

European Commission Directorate General for Energy (undated) DG ENER Working Paper The future role and challenges of Energy Storage, Brussels.

Evans, A, Strezov, V and Evans, T J (2012) Assessment of utility energy storage options for increased renewable energy, *Renewable and Sustainable Energy Review*s 16, 4141-4147.

Flatley, L, Giulietti, M, Grossi, L, Trujillo Baute, E and Waterson, M (2016a) "Analysing the potential economic value of energy storage", Institut d'Economia de Barcelona, Document de treball de l'IEB 2016/2.

Flatley, L, MacKay RS and Waterson, M (2016b), Optimal strategies for operating energy storage in an arbitrage market, *Journal of Dynamics and Games* 3(4), 371-398.

Forbes, K E and Zampelli, E M (2016), The accuracy of wind and solar energy forecasts and the prospects for improvement, Mimeo, Catholic University of America.

Greve, T and Pollitt, M G (2016) A VCG auction for electricity storage, EPRG working paper 1623 and Cambridge working paper in economics 1628.

Grossi, L, Heim, S and Waterson, M (2014) A vision of the European energy future? The impact of the German response to the Fukushima earthquake, University of Warwick, Department of Economics working papers, 1047.

Hutchinson, A (2015) Topics in Sustainable Energy: An Economic Analysis of Net Demand Volatility Management, University of Warwick, Unpublished PhD thesis.

Newbery, D (2016) Missing money and missing markets: Reliability, capacity auctions and interconnectors, *Energy Policy* 94, 401-410.

OFGEM (2014) *Capacity Market Rules*, London.

Zafirakis, D, Chalvatzis, K J, Baiocchi, G and Daskalakis, G, (2016), The value of arbitrage for energy storage: Evidence from European energy markets, *Applied Energy*, in press.

Table 1: Characteristics of wind forecasts and wind generation

	F1	FF	Wind gen
mean, MW	2972.6	3050.5	2631.6
sd	1829.5	1877.9	1650.4
min MW	225	207	71
max MW	7377	7450	6779
less than 1GW			18.89%

Key: F1 = initial forecast; FF = final forecast

Data from 7208 observations captured once per hour

Table 2: Regression of forward price on predicted load and forecast wind

IV regression of forward price (p) on forward load and forecast wind

	Coefficient	st error
Load	$7.41*10^{-4}$	$1.87*10^{-5}$

Wind	$-1.45*10^{-3}$	$-5.92*10^{-5}$

Notes:

7290 observations. Rsq 0.304

Instruments for Load are day, time, squares and cubes of these, interactions between these.

The Rsq for the first stage is 0.735

We are able strongly to reject the presence of a unit root in Price, using the Augmented Dickey Fuller test with up to 25 lags. Similarly, we are able to reject a unit root in Wind.

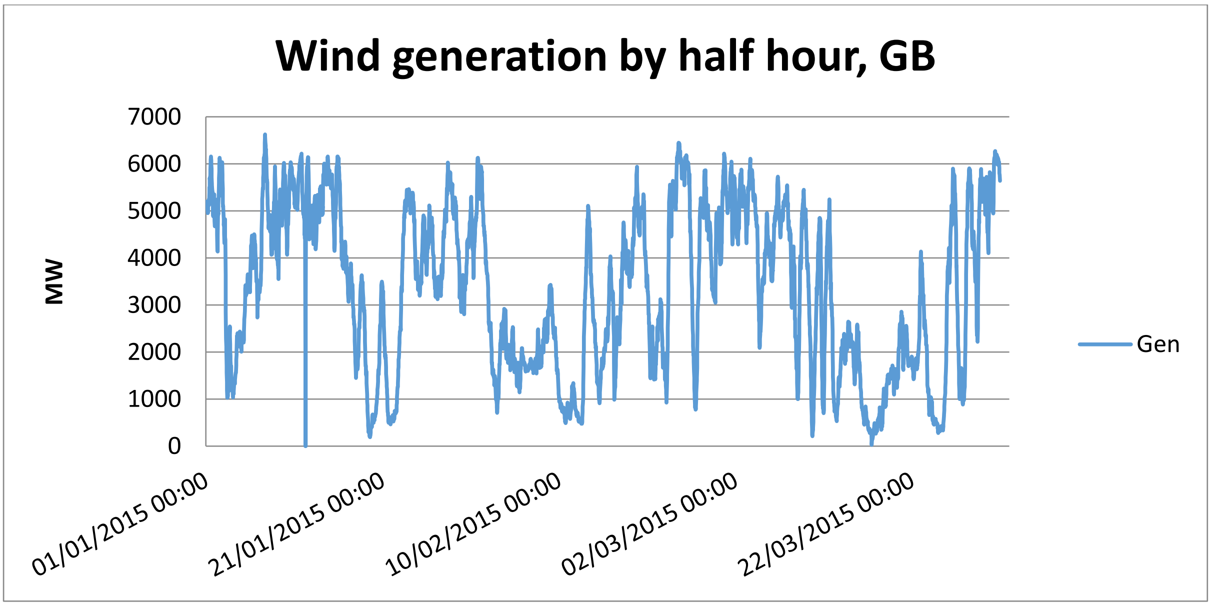

Figure 1: Three months' wind generation in Britain (source- Gridwatch)

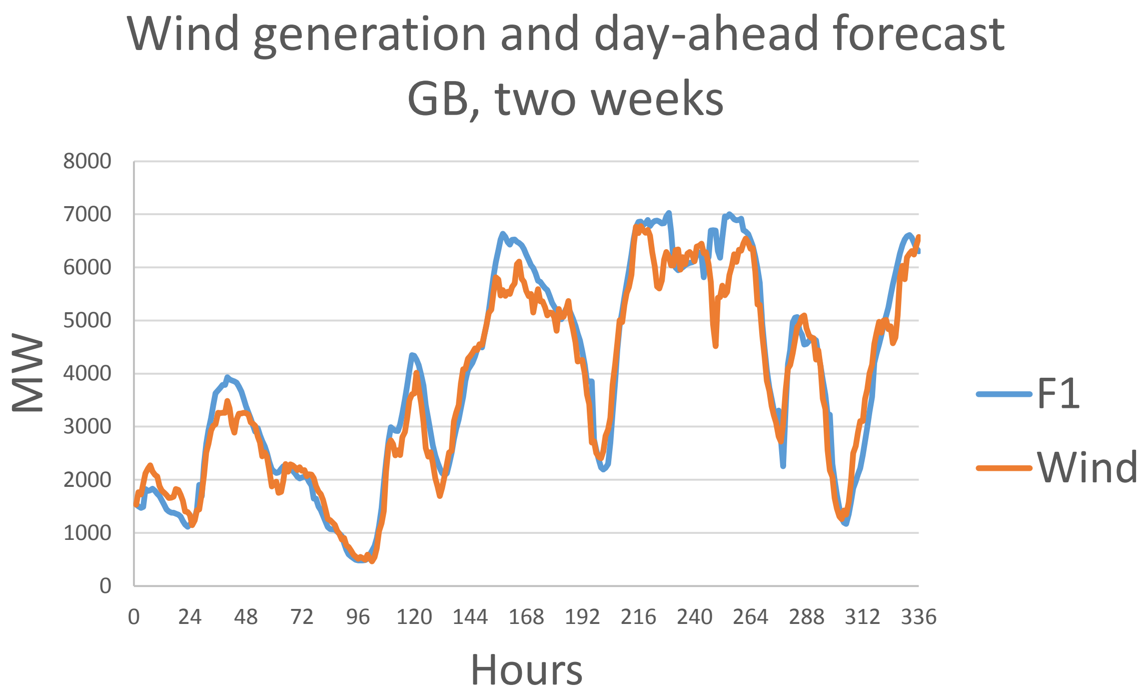

Figure 2: Short-term forecasting of wind in Britain (Gridwatch)

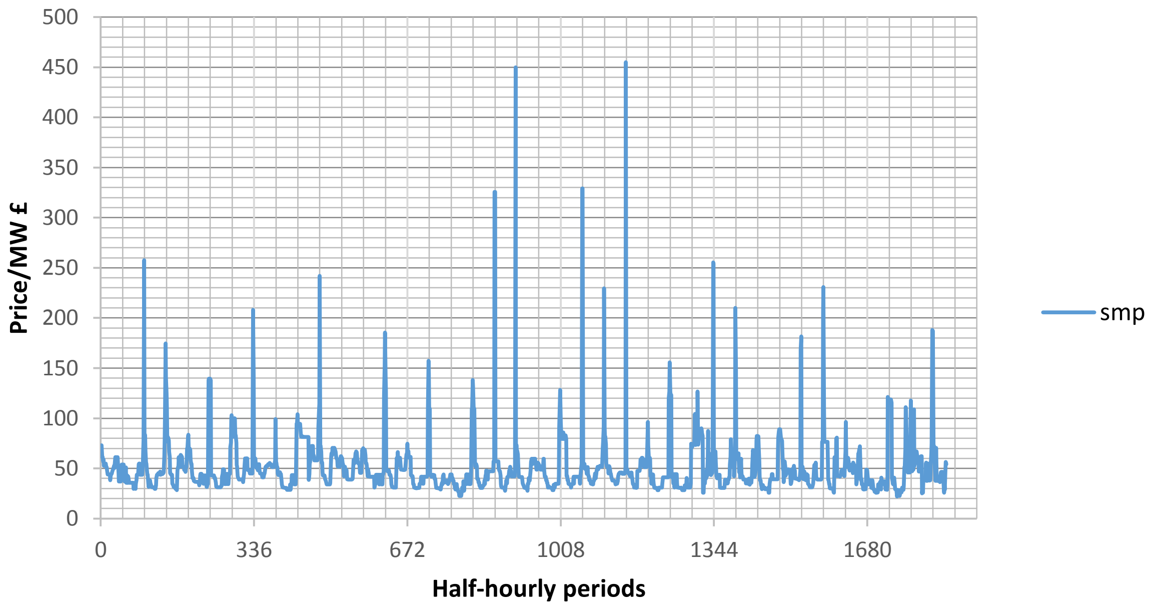

Figure 3: An illustrative of the pattern of half-hourly spot market prices in Britain (Elexon)

336 periods represents one week (7*48 half-hours)

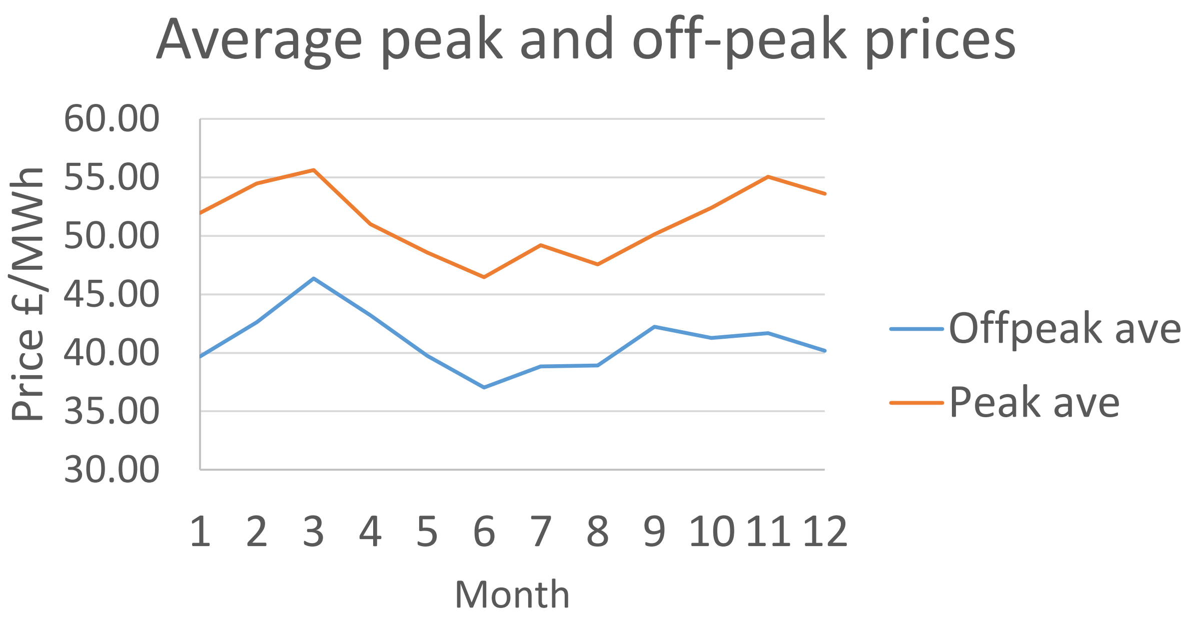

Figure 4: Monthly average wholesale electricity (over a three year period) (Elexon)

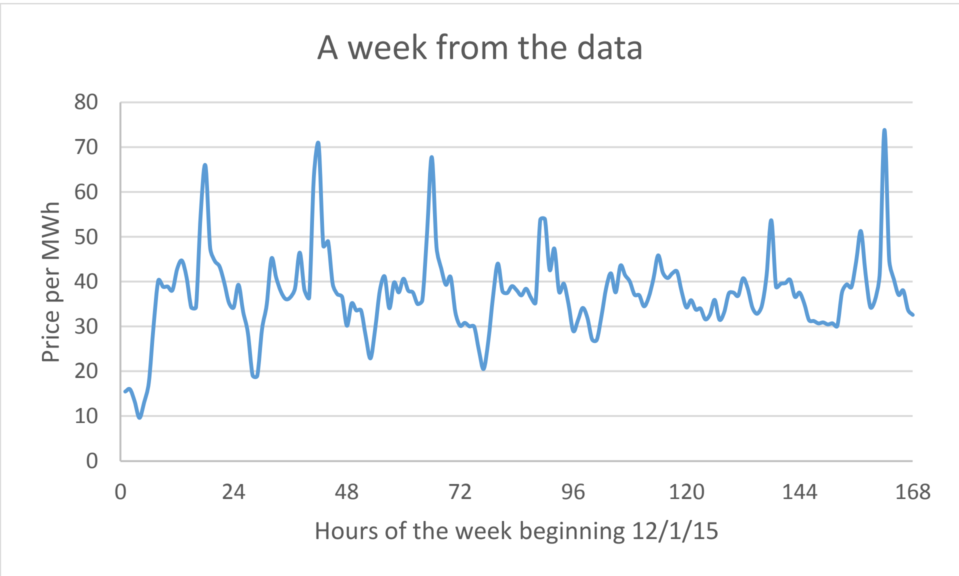

Figure 5: Illustrating the pattern of electricity prices over a week (Elexon)

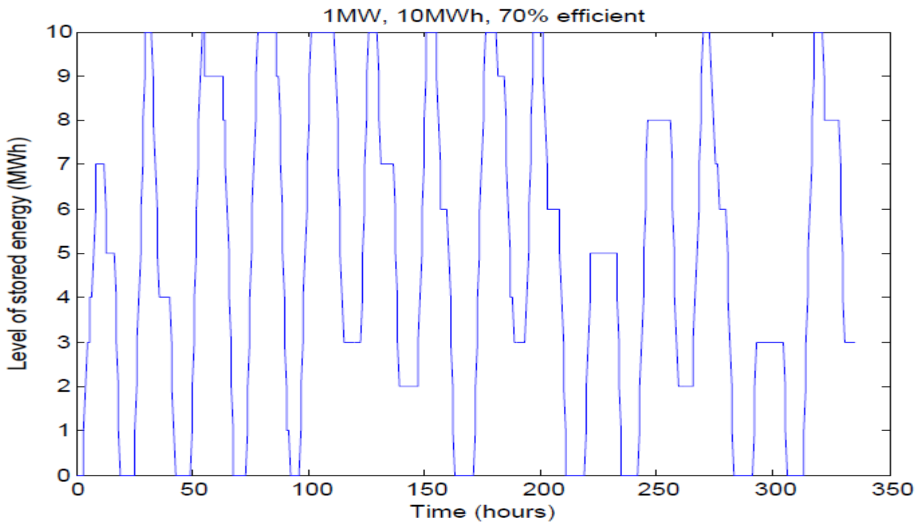

Figure 6a: Optimal operation of a small 70% efficient store over a two week period, assuming prices are known (Flatley et al, 2016b)

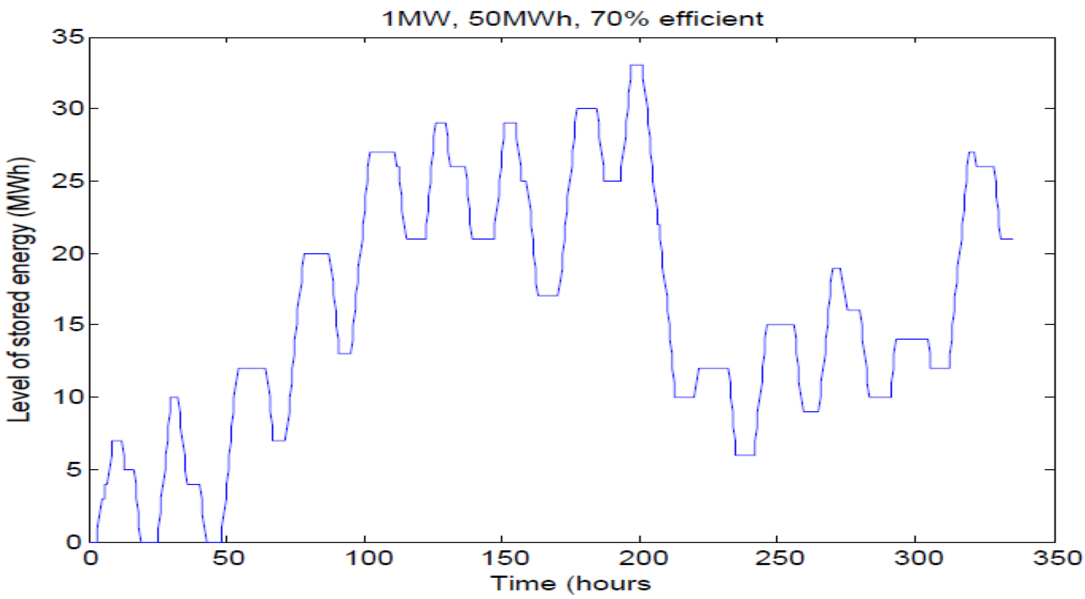

Figure 6b: Optimal operation of a 70% efficient but larger capacity plant over the same two week period. (Flatley et al, 2016b)

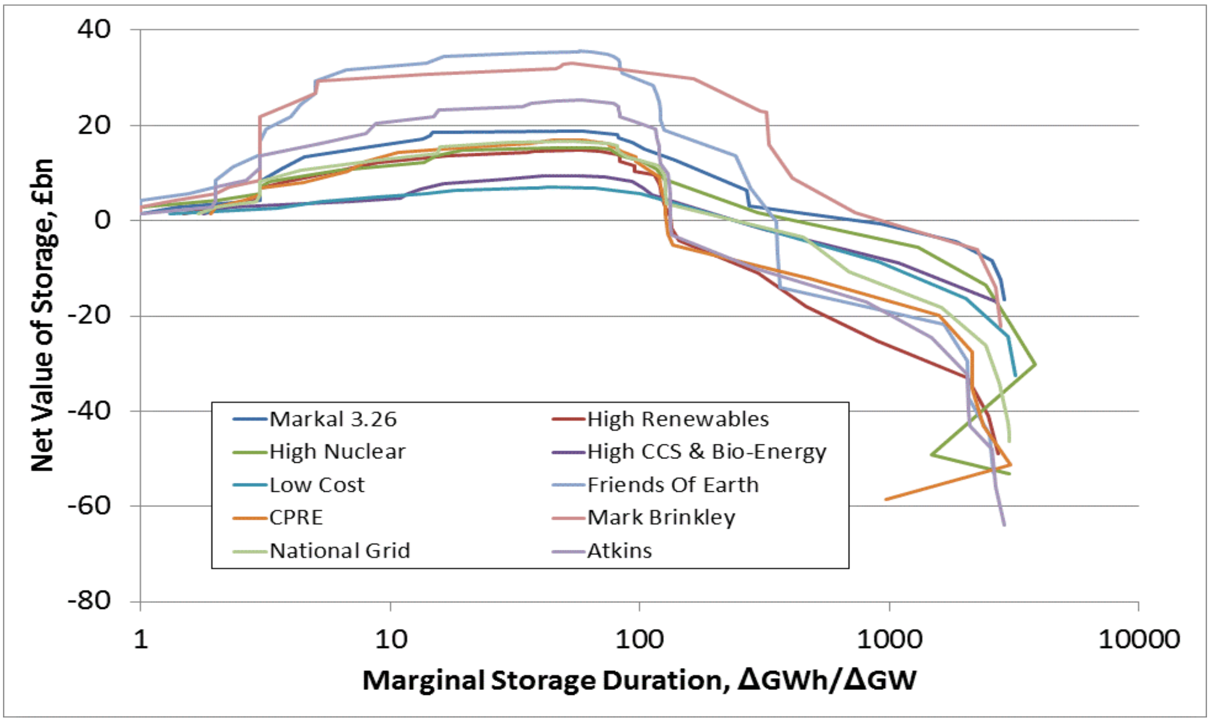

Figure 7: Net total lifetime value of storage on a cost-benefit basis under various demand scenarios, using engineering cost data. (Barton and Thomson, 2015)

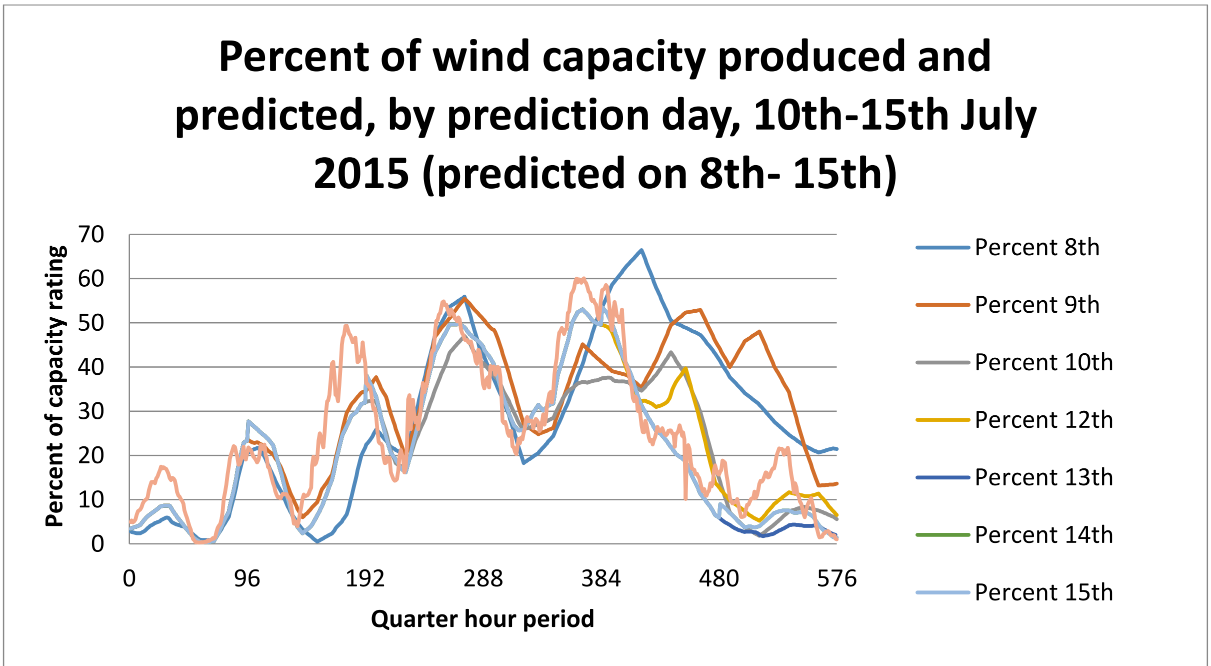

Figure 8: Wind generation predictions for Belgium over various look-ahead times for the same six days. (Own calculations from Elia data)