

Analyzing the Potential Economic Value of Energy Storage

Monica Giuliotti, * Luigi Grossi, ** Elisa Trujillo Baute, *** and Michael Waterson****

ABSTRACT

This paper examines the commercial opportunities for electrical energy storage, taking market prices as given and determining the extent to which a strategy of arbitrage across the day, buying at the lowest price times at night and selling at the highest price times during the early evening, and relying on price forecasts one day-ahead generates profits in the British context. The paper sets out the potential problems as the market moves to absorb increasing amounts of wind, then characterises the nature of prices, which reveals the importance of a strategy in which power is absorbed into store for a relatively few hours of the day and discharged over a relatively few hours. It argues that additional incentives may need to be put into place in order to render storage over relatively longer periods more attractive and to deliver broader social benefits which are unlikely to be generated and captured as a result of purely commercial considerations.

Keywords: Arbitrage; electricity markets; energy storage; market clearing; price forecasting.

<https://doi.org/10.5547/01956574.39.SI1.mgiu>

1. INTRODUCTION

Grid-scale electricity has two characteristics that distinguish it from every other good. Supply needs to equal demand at every point in time, to the second, or the system breaks down. Yet supply and demand naturally fluctuate over time. Storage to meet these fluctuations takes place via the conversion from other forms of energy, for example via combustion, or conversion of electrical energy into other forms of energy, for example from chemical energy in batteries, or releasing potential energy in the case of previously pumped hydroelectricity. In traditional electricity systems, development over time from local supply to national and even international supply systems has required system interconnection, which takes advantage of less than perfect correlation in demand across space, alongside precise monitoring and operation of flexible generating plant. Most of the burden of storage historically has taken place at the input level, via storage of coal, natural gas and, to a lesser extent, fuel oil. This approach to storage is however challenged by the increasing contribution of energy from renewable sources in modern energy systems.

The process of substituting renewable generating sources for fossil fuels in order to reduce carbon emissions has led to new and substantial challenges to the electrical systems of all countries.

* Corresponding author. School of Business and Economics, Loughborough University, LE11 3TU, UK. E-mail: M.Giuliotti@lboro.ac.uk

** Department of Economics, University of Verona, 37129 Verona, Italy. E-mail: Luigi.Grossi@univr.it.

*** University of Barcelona and Chair of Energy Sustainability - Barcelona Institute of Economics, 08034 Barcelona, Spain. E-mail: Elisa.Trujillo@ub.edu

**** Department of Economics, University of Warwick, Coventry CV4 7AL, UK. E-mail: Michael.watson@warwick.ac.uk:

These are not widely appreciated by the green lobby, the public more generally, or economists.¹ Yet they are of vital importance to the maintenance of readily available electricity supply which is a cornerstone of economic activity as we know it. Most renewable resources, for example wind and solar power, are intermittent rather than firm and are not biddable in response to price. Enhanced interconnection provides a partial solution as do demand management and storage. The traditional method of storage is via pumped hydro facilities, but these require geological features that are not easily replicated. Thus, alternative approaches that are scalable involving compressed air, heat or battery storage have been raised as possible solutions.

The main purpose of our paper is to examine whether there are sufficient economic incentives in place for such storage to operate profitably by exploiting intra-day arbitrage opportunities in the British wholesale market. Given the relatively small size of existing storage facilities in Great Britain compared to the size of the electricity system it would be unrealistic to presume that storage facilities currently operating in the wholesale market would be able to affect the market price much (see Renewable Energy Association, 2017). For this reason the paper focuses on the identification of an optimal intra-day strategy for operating a price-taking storage facility.

Several contributions in the literature on energy technologies have considered the issue of energy storage within a broadly economic framework, commonly with a particular institutional background in mind. Examples which rely on an explicitly economic framework include Walawalkar et al. (2007), Sioshansi et al. (2009), Sioshansi (2010), Carson and Novan (2013), and Bradbury et al. (2014) who attempt to assess private and social benefits from storage in the US context, while Zafirakis et al. (2016) investigate arbitrage opportunities in European markets. These papers address a range of issues, but one of the key aims is to analyse the impact of storage on prices; another is to evaluate the profitability of storage. Recent contributions which have focussed on the profitability of large-scale storage facilities include Barbry et al. (2017), Cruise et al. (2016) and Nasrolahpour, et al. (2016). These studies investigate the strategic choices available to storage operators who are able to affect the wholesale market price as a result of the large size of their facilities. Our work focuses instead on price-taking facilities, a more realistic assumption in countries where storage activities are at a relatively early stage of development and subject to regulatory uncertainty, as in the UK.

There are broadly two economic approaches that can be taken in evaluating the optimal level of storage in an electrical network (in particular in the context of significant intermittent generation). One is to evaluate the social benefits and costs of storage; the other is to examine the private benefits and costs of its operation. Storage activities which facilitate the balancing of supply and demand in energy system with high penetration of non-biddable generation produce a range of social benefits which could include:

- Saving on the capital expenditure required to install peaking plants (although these costs should be compared to the capital costs for the construction of storage facilities).
- Reduced expenditure on transmission and distribution grid reinforcement.
- Avoiding curtailment of output from renewable energy generation.
- Fuel saved through reduced ramp rates.
- Reduced need for low efficiency plants to operate.

However, not all these potential benefits can be captured by arbitrage, or through the provision of ancillary services on the basis of contractual arrangements, since the benefits arising from

1. This is partly because engineers (initially) and economists have been so successful in devising systems to deal with the problems, meaning blackouts are relatively rare. However we also know that, once a widespread blackout occurs, retrieving the system proves time-consuming and difficult.

these services are spread widely across different industrial participants with limited ability to appropriate the full rent from the socially beneficial activities they undertake. More succinctly there is a missing market problem² due to uncaptured positive externalities from storage activities. At the same time, some of the balancing activities potentially undertaken by storage facilities represent a form of business stealing as they would replace activities previously undertaken, and rents previously captured, by high cost peaking plants.

The approach we take here is to focus on the private ongoing benefits and costs, on the assumption that storage will be operated commercially. In the longer-term, as conventional generation and distribution resources are retired or replaced, broader considerations come into play, but we do not discuss these here.³ Nevertheless, implicitly we are considering stores capable of holding hours' rather than minutes' worth of power. More precisely our simulation results are based on technical parameters and operating costs associated with Compressed Air Energy Storage (CAES) technology, as published in Lund et al. (2009) and on the plant and pump size of the Hunthorf CAES plant in Germany, this first CAES plant in the world and the only one currently operating in Europe.

We take an optimistic stance to examining the commercial benefits of storage. All we ask is whether the first commercially-operated store to enter the market, operating through arbitrage, more than meets its running costs, ignoring the costs of construction, which will be site and technology-specific and may qualify for subsidy. Indeed, the cost estimates for power and capacity related costs for CAES technology vary widely depending on technical characteristics (e.g. adiabatic or non-adiabatic) of the plant, location and conditions of the cave used for storing compressed air.⁴ Because the store we consider is too small to influence prices, or generators, we take existing market prices as a good representation of what the store will face. In practice, the first store may be followed by others, stores may interact to reduce returns for other stores due to the effects on prices; generators may modify their strategies in response to the introduction of stores, etc. (see Hutchinson, 2015 chapter 2, and Cruise et al., 2015, for explorations of this issue). But if the first store cannot meet its running costs then such storage is infeasible as a commercial technology on the basis of arbitrage. Besides considerations about expected revenues and running costs, the operation of first store will also be affected by risk assessment considerations, as we discuss in relation to our simulations results. There are other potential revenue streams, as well as arbitrage, as we discuss later, but none clearly available for storage at present.

The questions we examine are: (i) whether there is a clear requirement for storage, (ii) what will be the focus of commercially-provided storage and (iii) whether this focus appears sufficient to tackle the issues the system faces. Our plan in this paper is to explain the requirement through considering briefly the characteristics of wind power, the prime source of renewable energy in Britain. We explain why interconnection and demand management, as currently operated and planned, are

2. Newbery (2016) discusses “missing market” issues in the context of the wholesale electricity markets, by explaining that ‘investors need confidence that the revenue they earn from the energy markets (...) will be adequate to cover investment and operating costs. If this revenue is (...) potentially adequate, but not perceived to be so by generation, companies or their financiers, then there is a “missing market” problem’.

3. An interesting and careful analysis that bears upon the social impact of storage is Gowrisankaran et al. (2016) on the social impact of solar power at various levels in the Tucson, Arizona district. However, this does not explicitly consider storage, including it (“perfect storage”) only via the beneficial impact on smoothing output of solar power.

4. Power related costs have been estimated to be in the area of 960–1250 \$/kW and capacity related ones between 60 and 150 \$/KWh (in 2010 US dollars terms) by ECOFYS (2014) for the case of underground CAES technology. A previous report by Stewart et al. (2009) also contains wide ranging estimates for capacity related costs between \$1.75 and \$40/MWh. These would need to be compared with our estimated profits of about \$4.8/KWh yearly profits in 2010 dollars, consistent with the basis of the ECOFYS figures.

a limited solution. We then move on to examining the characteristics of prices, which determine the nature of likely arbitrage activity, so in turn the likely characteristics of a store. Stores of different efficiencies and sizes are discussed. Finally, we move on to considering whether foreseeable future changes in the market will have an influence.

In order to provide evidence about the relatively novel and under-researched issues discussed above it is necessary to rely on a range of different sources of data with different time lengths. The analysis of arbitrage opportunities for a 'generic' storage facility (discussed in Section 3) requires historical data over a reasonably long period of time in order to see how well prices can be predicted and used effectively for the development of an optimal strategy to operate the storage facility profitably. On the other hand, the investigation about the ability to use weather forecasts at different horizons to predict the expected wind output and its impact on prices (discussed in Section 4) requires the use of data on actual and forecast wind outputs which is obtained from different sources from the price data and is only available for a more recent period of time (starting in November 2014).

2. THE REQUIREMENT FOR STORAGE


Storage is desirable if (a) there is significant variability in the source of power which requires some tailoring to conform to the desired profile of power demand and (b) alternative approaches are insufficient (or insufficiently competitive) to satisfy this tailoring requirement. We argue that both these conditions are satisfied in the case of Britain.

Wind power has important characteristics that provide challenges to modern electrical power systems. It is naturally unbiddable, it is extremely variable, its variability is time-correlated and it is difficult to forecast beyond a short period. Figure 1 below illustrates these features for Great Britain focusing on the year 2015, which will be used in our regression analysis and discussed in section 4. Mean generation in 2015 was 2671MW, with a large standard deviation of 1641. At times, virtually nothing is generated across the whole of Great Britain, at other times, almost 7 GW (maximum wind generation in 2015 was 6587 MW), or over 20% of average load, is generated using wind, and given current developments this will increase significantly.

There are periods within our data sample when one or the other happens for hours at a time, for example in 2015 for 18.8% of the time, generation is below 1GW, as illustrated in Figure 1. Time correlation is not easily seen in the figure, but in 2015 there were 22 occasions on which wind generation was less than 1GW for more than 24 hours. The longest such period was 74 hours, i.e. it covered just over 3 days, while across the whole sample the maximum period with wind generation below 1GW lasted over a week (luckily, in June). We also see that the day-ahead forecast, plotted in this graph, is highly correlated with the actual value, though it over-predicts significantly, by over 19% on average. Moreover, once we move beyond a couple of days ahead, predictions become extremely vague (the differences between actual and predicted values are very wide).⁵ This leads to potential problems for long term storage.

Demand management, for example through real-time pricing, has potential, but naturally cannot tackle the problem that consumers want the lights on when it is dark, more than when the sun is shining and generating solar electricity. Demand management shades into storage, because

5. This is an inference in the case of Britain; figures for Belgium, which also has a substantial proportion of its wind farms off-shore, suggest that predictions for five to seven days are rather uninformative. The problem is inherently difficult, since wind forecasts themselves are subject to increasing errors the further ahead they are made, and power generation is nonlinearly related to wind speed, but is zero below or above critical speeds.

Figure 1: The characteristics of wind generation (Year 2015)

storage can be implemented as either a response from users or an action taken by suppliers. On the part of users, immersion water heaters could (and some do already) heat water only when electricity is cheap, and those consumers with electrically-powered cars could charge up when electricity is cheap, and supply back to the grid when it is expensive. These incentives are similar to those involved in grid-scale storage, so although we focus on the latter, with some modification our analysis can take account of more consumer-based response in addition.⁶

Interconnectors respond to spatial differences resulting from different existing generation facilities and different consumer habits across space. These operations make it possible for interconnectors to obtain revenues exclusively through price arbitrage.⁷ By contrast, storage can respond to inter-temporal differences in loads resulting from habits across time (such as across hours of the day) which are substantial and, something of increasing importance, differences in supply availability from renewable generation. Consumption is indeed imperfectly correlated across countries, but there are many hours in the evening when the lights are on all over Western Europe, and many hours in the day in which all shops, offices and factories are operating. There is substantial interconnection already between the electrical systems of continental European countries. This assists in smoothing net loads across countries, so that for example France, with its large and relatively inflexible nuclear fleet, can import electricity when domestic demand is highest and export when it is lowest. Most such interconnectors are limited in capacity, manifesting in continuing price differences across borders.

Due to its island nature, Britain is less connected electrically with other countries than is typical for West European countries. Nevertheless, Britain has interconnection with France and the Netherlands (plus two with the island of Ireland), and further interconnectors are planned, the first two of which may come on-stream in 2019. These function as commercial operations on the basis

6. More significant modifications to the analysis are required to include the effect of electrical vehicle response.

7. A slight caveat: interconnection with Norway allows storage across time, since advantage can be taken of the substantial pumped hydroelectric reserves available there.

of arbitrage on prices, due to differing marginal generation and load balances between Britain and its near-continent neighbours. That is, to the extent to which different influences play a part in determining prices within Britain and its connected countries, the interconnector will operate to move electricity from the lower-priced country to the higher-priced. Of course these interconnectors experience line losses and also conversion losses, since there is a double conversion from 3-phase AC to DC (to facilitate undersea HVDC cable transmission) and back again at the other end. However in total such losses amount probably to less than 10% of the power transmitted.⁸ The total capacity of the French, Dutch, Norwegian and Belgian existing and planned interconnectors is around 6GW. Therefore, even if all are operating at full potential into Britain, they are only sufficient to tackle current wind fluctuations and are some long way short of being able to handle wind fluctuations once mean wind generation gets to around 20% of total power generated.⁹

At the same time, the storage requirements in a completely renewable system of operation are unfeasibly large for commercial operation. This is well-illustrated through a recent paper (Esteban et al., 2012), although this is not the point the paper makes. It calculates the required magnitude of storage in a 100% renewable system in Japan (with solar at 31% of the total capacity) as 41TWh! Moreover, as is clear from Figure 11 in Esteban et al.'s article, for much of the time this storage would be barely used; there is a long and almost continuous withdrawal from store in the period from July to September.

Commercial stores of natural gas are able to operate on the basis of an annual withdrawal and charge cycle, but most of the variance in gas prices comes over the long term in an annual cycle, whereas in electricity it comes over extremely short intervals, commonly a day, as we show below. Thus there are two dramatic differences in storage as between gas and electricity, one being the nature of price movements over time, the other the physical space required for storage. Rough, a depleted gas field in the North Sea that comprises Britain's main store of natural gas coincidentally holds up to 40TWh of stored gas, which would amount to less than half that in terms of generated electrical power, but the energy content of compressed natural gas per cubic metre is much higher than in electricity storage media such as compressed air, so if used for compressed air storage it would not amount to anywhere near the equivalent amount of electricity.

It is likely that both interconnection and storage will coexist as we move to a significantly less biddable system of generation, because solar power does not work in the evening in winter when demand is highest and wind power is inherently intermittent. Moreover, motivating a store to operate may be a more straightforward means of moving supply and demand into balance.


3. ARBITRAGE AS THE FOCUS OF COMMERCIAL STORAGE

In this paper we concentrate on the classic focus for electricity storage: arbitrage. Here we scope out the potential for additional storage in the UK energy system by analysing existing pricing patterns. Subsequently, we will consider whether the incentives are likely to change in the foreseeable future.

One of the characteristic features of wholesale electricity prices is a consistent gap between prices during peak hours and prices in the middle of the night (off-peak), arising because the

8. Entries in Wikipedia suggest losses from modern HVDC converters are approximately 2–3% at each end, whilst line losses are around 3.5% per 1000 km. (Pages on HVDC converters and HVDC, accessed 15/12/15).

9. Here we note that the island of Ireland, which already generates on average 20% of its power from wind, experiences very similar fluctuations to those exhibited in Figure 1. Hence it is unlikely that the extent of fluctuation is reduced much as a result of a denser mesh of wind generators in Britain in future.

Figure 2: First three months of 2014

marginal fuel technology for generating at peak has relatively high marginal cost. This provides an opportunity for electricity storage to act as arbitrage between these periods. Of course, since storage is less than 100% efficient, the variation between periods needs to be large enough that it more than covers the losses engendered by storing, which vary with the storage technology. What scope is there for arbitrage earnings to encourage operation of storage?


Here, we work with APX mid-prices (or market index data prices) obtained from Elexon¹⁰ over the four year period January 2011 to December 2015 to provide answers. These prices are used by Elexon to calculate the reverse price for each settlement period and ‘reflect the price of wholesale electricity in the short term market’ (<https://www.elexonportal.co.uk/article/view/188?cachebust=s-0bwnd5cy9>). They are calculated as a weighted average price for all qualifying contracts.¹¹

Figure 2 illustrates the pattern of half-hourly market prices over the three month period from January to March 2014.¹² The graph shows two things clearly. First, wholesale prices fluctuate a good deal, from just under £30 per MWh to well above the highest retail price. Second, the underlying pattern is diurnal. This second point is demonstrated more clearly, together with a third more subtle point in Figure 3, which takes a week (Monday to Sunday) from these data and examines it in more detail. Here we see that within each day, it is only a comparatively few hours for which the price is very high; this normally happens around 6 pm in the evening. The effect is slightly lowered at the weekend. The lowest prices occur at dead of night, between around 2 and 4 am. Between

10. Elexon is the balancing and settlement code company which manages electricity trading arrangements in England and Wales (<https://www.elexon.co.uk/>).

11. ‘A qualifying contract may take the form of any of the (...) relevant EPEX Products, specifically traded on or after the third calendar day prior to Gate Closure for the relevant Settlement Period.’ (see Market Index Definition Statement version 8, 23/2/2017, Appendix 1 https://www.elexon.co.uk/wp-content/uploads/2017/02/MIDS_v8.0.pdf).

12. We chose a three month period simply to facilitate the graphical presentation of the features.

Figure 3: Wholesale prices in the week starting 13 January 2014 (Monday)


November 2011 and November 2014 the average peak price was £51.26 per MWh, the average off-peak price was £40.94 per MWh. This influences the whole strategy that an arbitrating store would need to follow.

In order to make money simply buying off-peak and selling peak, a store would need to be at least 80% efficient. Early experiments over a period of several years that we have performed involving a risk free strategy of buying off-peak (or baseload) and selling at peak in a forward market (day/week/month ahead) is unlikely to expose sufficient gaps in prices (the “park spread”¹³) for arbitrage to be feasible for a store which is say 70% efficient. There are occasions where the gap is large enough, but they seldom occur (Giulietti, Grossi and Waterson, 2013).

However, as can be seen for the example day in Figure 4 (Monday 13th January, 2014), within the day there are certainly peak times when the price is over twice as much as it is in dead of night, so clearly there is potential for profitable diurnal arbitrage. Thus a more subtle purchase and sale strategy needs to be undertaken, in which particular periods showing the lowest prices and particular periods showing the highest prices, within the peak and off-peak periods, are used. This leads in turn to two considerations. The first is the appropriate ratio of store size to pump size, which will itself be dependent upon the round trip efficiency. The second is the development of a strategy that is able to look ahead in a relatively sophisticated manner to decide on when to take in and when to release energy from the store. Flatley et al. (2016) examine the nature of such an optimal strategy given the parameters of the store, focusing on the look-ahead time involved.

Returning to consideration of diurnal storage, clearly it is inappropriate to have a store capable of carrying a large amount of energy if the purpose is to charge and discharge over relatively few periods. If the store has pumps capable of filling the store with a power rating of P MW, and the

13. We define as a ‘park spread’ the profitability (in percentage terms) of storing power purchased at a given date t and sold on the same date via a contract signed at time t for delivery at time $t + N$: $\text{Future price}(t, t + N) - \text{day-ahead price}(t) / \text{day-ahead price}(t)$.

Figure 4: An example of a fairly typical daily pattern

operator envisages the store filling over a maximum period of three hours (and discharging over no greater a period), then the size of the store need be no larger than $3P$ MWh.

There are two main dimensions influencing the storage facility. The first is obviously the round-trip efficiency of the store, expressed as power that can be sold into the market divided by power taken from the market. The second, less obvious, is the size of the store relative to the size of the maximum output. However, the optimal size of store, given pump rating, is in fact endogenous, because it is influenced by the efficiency rating of the store.¹⁴ The more efficient the store, the longer is the range of time over which energy can profitably be discharged at peak (and charged off-peak). Thus, given market prices, there are two determinants of the profitability of the store: turnaround efficiency and the number of periods of intake and output, and together these impact on optimal store size. While a range of efficiency rating and pump ratings are technically compatible with stores of different sizes, the determination of the optimal size of a store will depend on its ability to at least recover the short term variable costs associated with the store operation, which in turn depends critically on the diurnal difference in wholesale prices allowing the store to exploit arbitrage opportunities.

We have seen that there are only a few periods in which daily prices are at their peak. This limits the nature of storage which earns revenue through arbitrage. Moreover, there are only small sets of consecutive times for which prices are very high at peak compared with prices at off-peak. Therefore, in order to assess the potential profitability of a storage plant operating on a daily basis by exploiting arbitrage opportunities in the balancing market, we performed a series of simulations aimed at identifying the optimal number of periods (measured as half-hours to reflect the bidding frequency in wholesale markets) when the store should operate in order to maximise profits through

14. This implies that the ability of a storage facility to exploit price variation for arbitrage purposes will depend on the charge and discharge rates and hours which are allowed by the relationship between store and pump size. For our calculations we use actual dimension of the Huntorf plant, which imply a charge and discharge rate of 1 hour, as discussed in more detail in section 3.2.

arbitrage. The methodology underlying the profitability simulations and the key results of such simulations are presented in the sub-sections that follow.

3.1 Methodology for profitability simulations

Suppose the plant has efficiency rate (output energy to input energy) of $\theta (< 1)$. In addition, the plant has input costs of x per MWh and export costs of y per MWh. So, it has running costs of $[(1 + \theta) \cdot x + y]$ per MWh of electrical energy produced.

Suppose that the plant is either charging at the full rate, or discharging at the full rate, or is doing nothing, and that there is no leakage. It has an input pumping capability of α/θ MW and an export pumping capacity of α MW.

Therefore, if it operates by inputting for τ settlement periods (usually half-hours)¹⁵ in a day, and exporting for the same period of time, it will input a total of $t\alpha/\theta$ MWh and export a total of $t\alpha$ MWh. In order to do this, the plant's size (working capacity) will need to be $\tau\alpha / \theta = A$ MWh. It is assumed that the size of the plant is small enough that it does not influence market prices.

What determines the optimal value for τ ? It will not be worthwhile operating the plant if the costs exceed the benefits. The benefits can be represented as the difference between the price per MWh of the input electricity and the efficiency-weighted price per MWh of the export of electricity. That is, in producing 1 MWh of electricity export, the benefit will be: $[p_y - (1 + \theta) \cdot p_x]$. This will change every period, but the approach that recommends itself is to start with the lowest period for the input price p_x and the highest price export period p_y and then to move to the next lowest input price and the next highest export price, and so on. The half-hour settlement periods are divided in "peak" and "off-peak" periods. Because in peak hours prices are usually higher than in off-peak hours, we charge the plant during off-peak periods and discharge during peak hours. Then, all possible combinations of peak and off-peak periods are analysed. In producing A units of output (MWh) by pumping as assumed, the benefits will be:

$$\frac{A}{\tau} \cdot \sum_{t=1}^{\tau} [p_{yt}^d - (1 + \theta) \cdot p_{xt}^d] \quad (1)$$

where p_{yt}^d and p_{xt}^d are the export and import price, respectively, at day d .

The costs, in terms of pumping costs (assumed to be the only variable costs) are constant per MWh at:

$$A \cdot [(1 + \theta) \cdot x + y] \quad (2)$$

Subtracting (2) from (1), experimenting with different values of τ , starting with 1, moving to 2 etc., results in an optimum number of periods, where each period at least adds something in terms of marginal net benefit. Notice that only the benefits side incorporates τ .

Once the optimal τ is found, we get an expression in terms of A . This is:

$$A \cdot \left\{ \frac{1}{\tau^*} \cdot \sum_{t=1}^{\tau^*} [p_{yt}^d - (1 + \theta) \cdot p_{xt}^d] - [(1 + \theta) \cdot x + y] \right\} = A \cdot B \quad (3)$$

15. In our simulations we use prices at half-hourly frequency as this is current trading frequency in the UK wholesale market. However if real time prices were introduced these would increase the opportunities for a storage facility to exploit sharp price variations due to demand imbalances, within the limits of the technical characteristics of the plant, and therefore increase its profitability. We thank an anonymous reviewer for highlighting this important point.

Suppose the construction costs of a plant depend only on its size, i.e. they are $k \cdot A$. Then, the plant will be worthwhile building if the net benefits over time of building the plant are positive, that is

$$B(1 + r + r^2 + r^3 + \dots) \geq k \tag{4}$$

Notice that we need not know A to calculate this. One way we could evaluate this is to assume a discount rate r (e.g. $r = 0.95$) and hence evaluate a target construction cost per MWh of capacity.

In our simulations we aim to calculate the optimal number of periods when the storage facility should operate by evaluating the marginal benefit generated by increasing by one additional unit the number of periods when the store operates. In order to determine the optimal value of τ we devised a procedure in two stages which requires splitting the sample of half hourly electricity price data, covering the period from January 2011 to December 2015, into two parts. The period up to December 2014 (inclusive) is used to devise the optimal strategy for operating the store, i.e. in the first stage of the simulation the optimal buy-sell strategy is identified on the basis of historical daily data covering the interval $d = 1, \dots, D$.

In the second stage the optimal strategy obtained in the first stage is applied to daily data with $d = D + 1, \dots, D + h$, covering the period January 2015 to December 2015, in order to assess the difference between expected and actual profits when the optimal strategy identified in the first stage is applied to future time periods.

3.1.1 First stage

As discussed earlier, the half-hour settlement periods are divided in “peak” and “off-peak” periods. Taking τ as the number of periods needed to charge and discharge the plant each day, for the j -th combination of periods at day d the profit (excluding capital and O&M costs) is given by:

$$\pi_{dj} = (1/\tau) \sum_{t=1}^{\tau} [p_{yt} - (1 + \theta)p_{xt}] \text{ for } j=1, \dots, J \text{ and } d=1, \dots, D \tag{5}$$

For each combination of peak and off-peak prices, the results can be summarized over the historical data ($d = 1, \dots, D$) using different location and scale indexes. The most obvious is the arithmetic mean $M(\cdot)$, that is:

$$M(\Pi_j) = (1/D) \sum_{d=1}^D \pi_d \tag{6}$$

but in order to reduce the impact of outliers we can replace the mean with the median.

If we choose the mean of profits as the summary index, the optimal combination of peak and off-peak periods is the one which maximizes the chosen index over all J combinations, that is:

$$OPT = \underset{j}{\max} [M(\Pi_j)] \tag{7}$$

The same principle for selecting the optimal combination could be applied to other summary indexes such as the median or the ratio of mean to standard deviation. As the intraday structure of prices changes between weekdays and weekends, and also between winter and summer months, we have found different optimal combinations for weekdays in summer, weekends in summer, weekdays in winter and weekends in winter.

3.1.2 Second stage

In this stage the optimal buy-sell strategies identified in the first stage are applied for $d = D + 1, \dots, D + h$ in order to obtain an estimated daily profit, that is:

$$\hat{\pi}_d = (1/\tau) \sum_{t=1}^{\tau} [\hat{p}_{yt}^d - (1+\theta) \hat{p}_{xt}^d] - [(1+\theta)x + y] \text{ for } d=D+1, \dots, D+h \quad (8)$$

where \hat{p}_{xt}^d and \hat{p}_{yt}^d are the optimal buy and sell prices on day d .

To summarize the output of the strategy for the period $[(D + 1), \dots, (D + h)]$, we use different measures:

- average daily profit

$$M(\hat{\Pi}) = (1/h) \sum_{d=D+1}^{D+h} \hat{\pi}_d$$

- fraction of positive daily profits

$$FP(\hat{\Pi}) = (1/h) \sum_{d=D+1}^{D+h} I(\hat{\pi}_d > 0)$$

where $I(\cdot)$ is an indicator function assuming value 1 when the condition in bracket is satisfied and 0 otherwise.

- average value of positive daily profits

$$MP(\hat{\Pi}) = \frac{\sum_{d=D+1}^{D+h} I(\hat{\pi}_d > 0) \hat{\pi}_d}{\sum_{d=D+1}^{D+h} I(\hat{\pi}_d > 0)}$$

- standard deviation of daily profits

$$SD(\hat{\Pi}) = \sqrt{(1/h) \sum_{d=D+1}^{D+h} [(\hat{\pi}_d - M(\hat{\Pi}))^2]}$$

- standard deviation of positive daily profits

$$SDP(\hat{\Pi}) = \sqrt{\frac{\sum_{d=D+1}^{D+h} [(\hat{\pi}_d - MP(\hat{\Pi}))^2 I(\hat{\pi}_d > 0)]}{\sum_{d=D+1}^{D+h} I(\hat{\pi}_d > 0)}}$$

3.1.3 Stochastic procedure to find the optimal τ

In addition to the analytical approach to the search for the optimal τ discussed above, we also implemented a stochastic approach. The main difference from the analytical procedure is that one-day-ahead forecasts of half-hourly prices are obtained through an estimated time series model based on a rolling window procedure for each day of 2015.¹⁶ In particular, data from the beginning

16. The size of the rolling window is kept constant by adding a new observation to the original series of days and removing the oldest observation at each successive stage.

of 2013 to the end of 2014 were used to estimate SETAR(7,1) models¹⁷ with daily dummy variables for each of the 48 settlement periods.

Two different SETAR models have been considered: a classical estimator version (based on recursive OLS) and an outlier-robust version based on a polynomial weighting function (Grossi and Nan, 2015).

Once the predicted price values are obtained, the optimal buy-sell strategy is decided each day using the price forecasts for the next day. For instance, on 31/12/2014 we obtain the forecast prices for each settlement period of 1/1/2015 and, assuming $\tau = 1$, we decide the optimal buy-sell strategy buying at the minimum forecast price and selling at the maximum forecast price. Of course when $\tau > 1$, maximum and minimum are computed as average values within the period defined by τ .

3.2 Simulation results

The procedures described in section 3.1 have been applied to find the optimal τ under three realistic values of the efficiency rate for storage technologies (θ): 0.7, 0.8, 0.9. The first stage was implemented to define the optimal buy-sell strategy (optimal period to buy off-peak and to sell at peak time) using price data for the period 2011–2014, while in the second stage it was applied to price data for the whole year 2015. Our simulation results are generated on the basis of technical parameters typically associated with Compressed Air Energy Storage (CAES) technology.

Here, it is taken as given that the plant has been constructed, so the major part of the operating costs, the costs we take into account, are those involving charging and discharging the plant. The costs of pumping and turbining for a compressed air energy storage plant are given by Lund (2009) as 2.3€ per MWh and 2.7€ per MWh, respectively. We translate these values into £/MWh at an exchange rate of 1.25€ per £. Energy input is assumed to be purchased (spot) at mid-market Elexon prices and sold also at these prices obtaining when sales take place. The plant operates for a number of periods filling (dependent on size of pump relative to capacity) and the same number of periods discharging later in the day or on the next day. In our simulations of the scenario analysis we assume operating with a standard compressed air storage plant with a storage capacity of 580 MWh and a pump size of 290 MW.¹⁸ This means that the average charging/discharging time is one hour, thus we have to exclude combinations of charge/discharge with less than one hour interval between them. Furthermore, the price varies in the two settlement periods used for charging/discharging procedure, thus we have to take the average price of two contiguous periods.


Assessing an arbitrage procedure usually implies the computation of risk measures which must be used to correctly evaluate the return of the investment (Vatani et al., 2018; Nasrolahpour et al., 2016). For this reason, in the following figures the Sharpe Ratio (= Average Profit/Standard Deviation of Profits) will be shown alongside the average profits. The calculated indexes were obtained in stage two of the procedure for different values of θ (efficiency rate) and τ (number of periods of operation) using the arithmetic mean as the summary index to determine the optimal strategy in the first stage of the procedure.

The optimal value of τ was also determined in stage two by considering only positive profits and the proportion of positive profits across all values of τ . When different summary indexes were used in the analysis they generated similar optimal values of τ to the ones reported in the figures

17. SETAR models are autoregressive models where autoregressive parameters vary according to a threshold which is estimated simultaneously to autoregressive coefficients.

18. These parameters reflect the technical characteristics of the only CAES plant currently operating in Europe: Kraftwerk's Huntorf plant located in Germany.

Figure 5: Average profit and Sharpe Ratio (2015) by value of τ . Results rely on the optimal strategy obtained on the period 2011–2014.



below.¹⁹ For this reason, our discussion will focus on the results based on the mean profits as the main summary index. Our simulations seek to maximise average profits rather than the maximization of the total profit on a daily basis because this strategy would require a constant readjustment of operations in response to observed prices which is limited by the technical characteristics of the CAES plant.²⁰

Figure 5 (left panel) reports the value of the average estimated profits for the whole of 2015 associated with different values of τ . It can be seen for any value of θ the optimal value of τ , in this case is 2, with average profits increasing with turnaround efficiency. The Sharpe Ratio (right panel) confirms, as expected, the best performance linked to higher values of efficiency. However, the trajectory related to the highest level of efficiency (0.9) is flatter than the others, meaning that, for high levels of efficiency, when τ increases the level of risk decreases more quickly than the profit. Furthermore, for the highest level of efficiency, the optimality in terms of right balance between profit and risk is indifferently reached when τ is equal to 2 or 3.

Looking at the different plots for combinations of seasons and days of the week, it is interesting to note that the maximum average profits are obtained during weekdays in winter (Figure 6, left panel) while during weekdays in summer average profits are non-negative only for the case of θ larger than 0.8 (Figure 7, left panel). Furthermore, during weekdays in winter relatively high levels of profit are compatible with a range of values of τ between 1 and 4. It is very interesting to note that according to the Sharpe Ratio (Figure 6, right panel), the optimal value of τ in weekdays of winter 2015 is 4 and the shape of the trajectories are completely different than in all the other cases because the performance index clearly increases for values of τ from 1 to 4.

Therefore, combining an average profit index with a risk measure index enables to clearly identify the optimal value of τ . In weekdays of summer 2015 the risk measure follows a trajectory very similar to the average profit, thus the plot of the Sharpe Ratio (Figure 7, right panel) confirms the results displayed in the left panel of the same figure.

The results based on the stochastic approach, which rely on one-day-ahead price forecasts, are perfectly comparable with the analytical procedure because the selected optimal strategy is always within one day.

19. These results are not reported for brevity but are available from the authors on request. In particular, the output of the procedure when the summary index is given by ratio of standard deviation (risk measure) to the mean profit is very similar to that obtained when considering just the mean profit.

20. We thank an anonymous reviewer for raising this important conceptual point.

Figure 6: Average profit and Sharpe Ratio (2015, week days winter) by value of τ . Results rely on the optimal strategy obtained on the period 2011–2014.


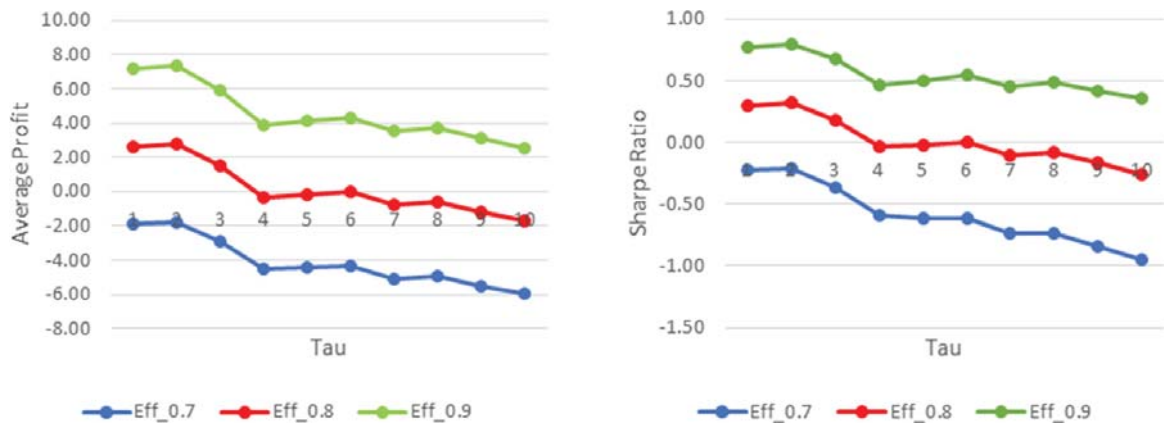


Figure 7: Average profit and Sharpe Ratio (2015, week days summer) by value of τ . Results rely on the optimal strategy obtained on the period 2011–2014.



Once predicted values are obtained, the optimal buy-sell strategy is decided each day using the forecasted prices of next day. For instance, at time T the forecasted prices for each settlement period of time $T + 1$ are available and, assuming $\tau = 1$, the optimal buy-sell strategy is to buy at the minimum forecasted price and sell at the maximum forecasted price within day $T + 1$. Of course when $\tau > 1$, maximum and minimum are computed as average values in the period defined by τ .


To give an idea of the main results obtained using the stochastic approach, the average profits and the Sharpe Ratios are reported in Figure 8 using a robust²¹ SETAR(7,1) model for the whole year 2015.

Compared to the results of the analytical approach illustrated in Figure 5 we find that $\tau = 2$ could also be considered as the optimal value in order to maximise average profits. Moreover, the average profits obtained with the stochastic approach for values of τ less than 5 are higher than those obtained using the analytical approach. Looking at the Sharpe Ratio (Figure 8, right panel), the trajectories are very similar to those observed in Figure 5, at least for the highest level of efficiency: the line is very flat for low values of τ and the best performance is even compatible with $\tau = 6$.

When performing similar simulations to those presented in Figures 6 and 7, which distinguish between days of the week and seasons, but adopting the rolling window approach discussed above, rather than applying the optimal strategy throughout the forecast period (2015), we find that

21. The outlier-robust SETAR model is estimated applying a polynomial weighting function (see Grossi and Nan, 2015).

Figure 8: Average profit and Sharpe Ratio (2015) by value of τ . Results rely on the day-ahead predicted prices obtained through the estimation of a robust SETAR(7,1) model.



the results for weekdays in winter are virtually unchanged, while for weekdays in summer the optimal value of τ increases to 3 for all levels of efficiency, although with similar estimated average profits for the forecast period (as illustrated in Figure 9, left panel). When the Sharpe Ratio is considered (Figure 9, right panel), the optimal value of τ is 4 when the efficiency is 0.9, while it is still 3 for lower efficiency levels.

In summary the results of our simulations indicate that average private profits can be positive for a sufficiently efficient store and can be maximised by charging and discharging energy over a relatively limited number of periods. The computation of risk measures combined with the average profits leads to conclusions which, in some cases, could differ from those obtained when just relying on pure profits. Our results also indicate that different strategies might need to be employed in order to maximise profits depending on whether the store operates during weekdays or weekends, and during different seasons.

4. FORESEEABLE CHANGES IN THE PRICE STRUCTURE

So far we have examined the possibilities for arbitrage under a number of limiting assumptions. We have assumed that the average gap between low and high prices is likely to remain constant into the future, although there are forces that will push this in both directions. Therefore we should consider the nature of these forces to some extent.²²

For example, increased use of “smart” technologies within the home and business may smooth power demand over the day relative to the situation at present, lowering margins between peak and off-peak prices. On the other hand, the German experience (Grossi et al, 2017) illustrates that a key factor relates to what is assumed about plant that is retired, once renewables become more important. If inefficient peaking plant is retired, the gap could grow, whereas if substantial baseload plant is retired, the gap may again narrow. But the answer actually depends quite subtly on the nature of the existing plant’s marginal costs relative to the marginal costs of the plant that is retained and it is not possible to give a general answer on this point. The benefits of arbitrage may either increase or decrease over time.

22. We maintain the assumption that the store is too small itself to influence prices significantly, although we are not specific on the size of the store. To put this point slightly differently, we assume there is sufficient capacity that can be purchased in the night-time period without causing an increase in prices, and (a lesser problem) the market is willing to take power in peak periods without the additional supply lowering prices measurably.

Figure 9: Average profit and Sharpe Ratio (2015) by value of τ . Results rely on the optimal strategy obtained through a rolling window procedure.




Figure 10: Illustrative movements in the merit order and their effect on average peak and off-peak prices

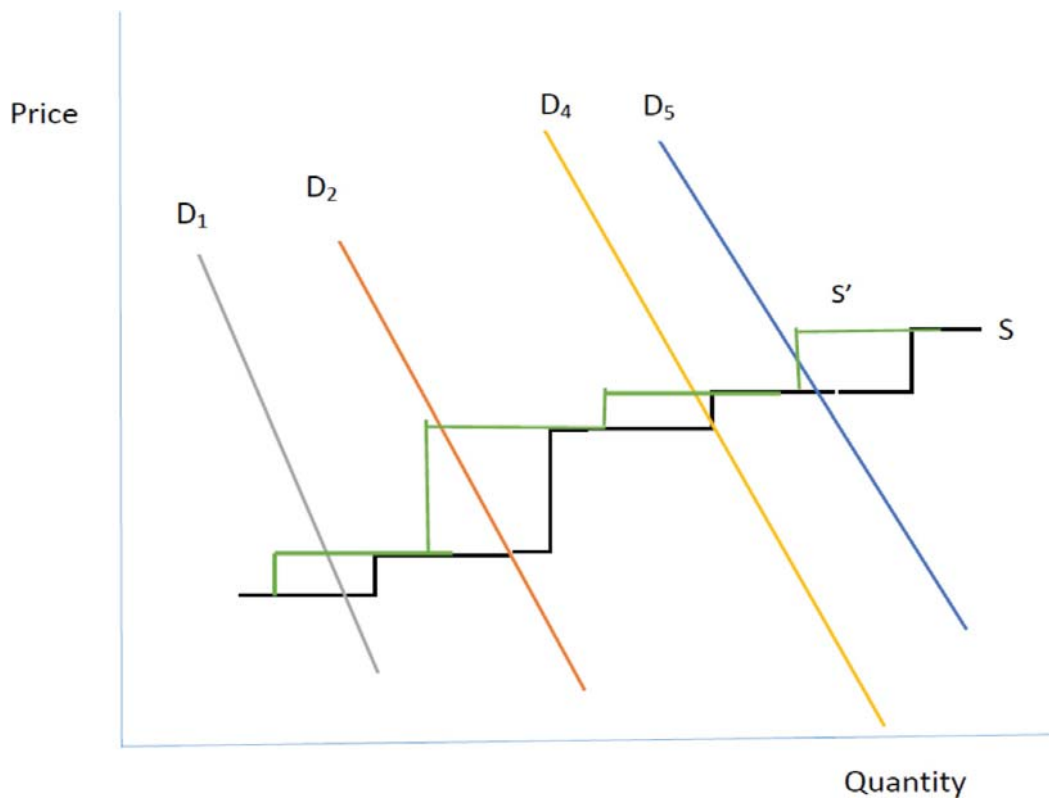


Figure 10 illustrates the ideas. Suppose initially that demand ranges between D_1 and D_2 off-peak, and D_4 and D_5 at peak. If some plant at the lowest end of the merit order or supply curve S is retired (say, nuclear powered and coal powered plant), then the supply curve shifts to S' , the green line. Both ranges have increased, but the lower range has increased more and the mean value of off-peak price has increased. However, this is due to the big jump in marginal cost (and so price) at the lower end of the merit order. The opposite effect could as easily happen if there is a big jump in the distribution towards the upper end of the merit order.

Will the undoubted growth in renewable generation help or harm a storage unit's profitability? To examine this question, we focus on the characteristics of wind generation in 2015 and 2016.

Table 1: Characteristics of wind generation

		Spring	Summer	Autumn	Winter
Variation from annual mean, percent	2015	-2.8	-25.9	-9.4	38.1
	2016	-12.0	-34.1	9.5	36.6
Time of day comparison		Percent Peak over average		Percent Peak over quiet time	
Means	2015	-5.5		-3.0	
	2016	6.7		9.1	

We first examine whether wind generation is lower or higher in off-peak times than at peak times, and more specifically whether it is lower or higher during the times when a store is most likely to be taking in electricity.²³

Table 1 contains data which bears on these issues. The overall, and somewhat unsatisfactory, impression from the data is that the vagaries of wind are such as to make any answer problematic. Looking at 2016, generation is higher at the peak than the off-peak, also higher than at the lowest demand period of the day. However, 2015 data yields the opposite conclusion! Moreover, total wind generation in 2015 is actually higher than in 2016, so we cannot infer that the difference in 2016 is that more has come on stream. Hence, as more wind comes on stream over time, the price gap between the highest and lowest price times might either increase or decrease.

Solar generation is even more difficult to assess, since National Grid does not measure it centrally and Gridwatch has only collected the data, based on a different and possibly dubious framework, since late October 2016. Instead, it is implicit in the Load data. Clearly on average there is more solar generation in the middle of the day and if, as in Germany, there were a midday peak in load, this would not be helpful to storage prospects. However, there is no evidence of a midday peak in the UK and the highest consumption hours, commanding the highest prices, tend to be between around 17.00 and 19.00, well after the time when solar generation is important.

Finally, looking over the year, there is obviously more solar generation in summer, when load is lower, as are prices. Table 1 shows us that it is markedly windier in the winter, on average, than in other seasons, but again there are substantial differences as between 2015 and 2016. Power prices do not show any particularly marked seasonal variations though, in part because of planned maintenance of major plant in times of lower load.

In sum, evaluating the impact of growth in renewables on price relatives is subject to significant uncertainty. Nevertheless, it is clear that wind generation, and wind forecasts, do influence power prices.

Turning to this issue, we look at price predictions based on wind forecasts. We have perforce to ignore the impact of solar energy due to lack of data, as explained above. We are also assuming away unexpected divergences from load so that load is affected only by standard factors such as time of day, day of week, etc., and there is by assumption no demand management. Therefore, prices will be determined by the factors influencing load in the short term, when fuel prices are given, after taking account of variances in wind. This is the basis of the relatively simple regression results outlined below.

Wind generation forecasts in Britain are only published for the next day and the following day (actually, 36 hours ahead) by National Grid. Since end-November 2014, these forecasts have

23. This exercise relies on generation data sourced from the Gridwatch website. There are occasional gaps in these data series, but they are not numerous enough to have a material impact on the mean values discussed below.

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

	Linear		Log log	
	Coefficient	Standard error	Coefficient	Standard error
Predicted Load	7.73×10^{-4}	2.07×10^{-5}	0.610	0.0164
Forecast wind	(-1.70×10^{-3})	5.89×10^{-5}	(-0.106)	0.004
R squared	0.278		0.287	
Observations	8760		8760	

Notes: Instruments for Load are: day, time, squares and cubes of these, interactions between these. The R squared for the first stage is 0.683.

been archived by the “Gridwatch” website. The exploration here focuses on data forecasting over the calendar year 2015.²⁴

We wish to predict price based upon load and wind generation, in order to examine whether there is a predictable pattern. The implicit assumption is that, apart from wind, other generation is essentially controllable and will involve a portfolio combination of gas and coal generation varied according to relative 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:

$$p_{t+1,t}^f = g(L_{t+1,t}^f, W_{t+1,t}^f)$$

where the meaning of $x_{t+1,t}^f$ is the forecast at t of the value in $t + 1$ for a generic variable ‘ x ’, p is wholesale power price, L is load and W is wind generation. A similar model could be used for longer-term forecasts. Load is instrumented for these purposes by a deterministic polynomial model of the day and time of day. Wind generation forecasts come from the Gridwatch site and price is the APX day-ahead forward mid-price for that half hour. The IV regression is shown in Table 2. The coefficients for both regressors are very significant.

Based upon the linear regression, we find that the elasticity of price with respect to load, at the mean, is 0.696. The elasticity of price with respect to the wind forecast is just 0.114. 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 £2.79, where the average price is £39.67 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, a significant calm period does not 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 dwarfs the impact of wind on price.

This leaves open the question of the impact of greater wind penetration and possible longer-term forecasts. However, similar experiments on price based upon day-ahead forecasts for Ireland, where wind is at 20%, suggest only a modest impact on price.²⁵ The 36 hour forecast reported

24. The Gridwatch data are very much “as is”. Substantial cleaning is required in order to render them suitable for analysis. The reports on which we base our analysis are hourly on the half hour, where gaps in reporting have been smoothed. We concentrate on forecasts of wind and ignore solar power, because until very recently this appears not to be measured as an output at grid scale in Britain, rather it feeds into the distribution network and impacts negatively on load.

25. Less formal estimates for Ireland suggest that a one standard deviation upwards (downwards) in the wind forecast, at the mean, creates a little less than a 7% shift downwards (upwards) in the day-ahead price. This is on the basis of a three month period in winter 2014/15, but without correcting for potential endogeneity of load and using the EA1 (first ex ante) run of prices.

via Gridwatch has a similar, although somewhat lesser, impact on price. But what of longer-term forecasts that might predict calm or a windy period for several days?

Here, there is a potential problem from the viewpoint of capturing market-relevant information to engage companies to store for several days rather than holding diurnally. Longer-term forecasts for Britain appear not to be published, but the Belgian system operator publishes forecasts for up to a week ahead. Simple calculations we have performed suggest that these are rather inaccurate beyond around three days out. Similarly, a Spanish commercial website (Meteologica) produces forecasts for example for German regions, but again beyond a few days, the differences between actual and predicted values on these are extremely wide. At present levels of wind penetration relative to flexible biddable plant, a day's notice is enough for planning purposes. However, it remains possible that, if wind penetration should increase significantly or biddable flexible plant becomes short, a commercial imperative to develop better long term wind forecasts may arise.

5. CONCLUDING REMARKS

We have shown that, given the current differences in peak and off-peak power prices, it is possible to profitably operate a store over a limited number of periods during the day, if we ignore the store's initial construction costs. Moreover, our analysis of the impact of wind output forecast on power prices indicates that at the current time there are insufficient profitability incentives to store power over more than one day in order to help balance supply and demand over protracted periods of wind calm.

Are things likely to change much in future? The characteristics of wind imply that whilst on average it can generate significant amounts of electricity in countries such as Britain, there is a very substantial variance around the amount produced. Moreover, there is a time correlation in generation that means periods of calm tend to occur together, as do periods which are exceptionally favourable to wind generation (not the most windy days, when generation has to be curtailed, but days with significant wind). If the electricity generation system relies to a significant extent on wind, it needs to find methods of covering the periods when wind is not generating. In particular, periods longer than a few hours are ill-suited to demand management tactics, because switching a refrigerator or immersion heater off for several days, or refraining from charging your electric car, are unlikely to be practical propositions. Nor, to the extent that wind generation may be low across the European continent for several days under certain climate conditions experienced in winter, does interconnection provide a good solution. This type of operation, holding power over several days, would not appear profitable on the basis of present price patterns within storage facilities given current and projected efficiency rates and might be better addressed by gas plants constructed for the purpose, for example.

Two alternatives suggest themselves. The first is that pricing patterns will change sufficiently to make it worthwhile to store for longer periods on an arbitrage basis. The second is that new means of incentivising longer term storage may need to be contemplated as an alternative to high-emission diesel generators. The obvious solution here is for storage to enter the capacity market,²⁶ since storage in media such as compressed air energy storage is technically feasible for periods in excess of a week, and arguably also socially feasible according to calculations by Barton and Thomson (2015). At present the incentives to encourage this appear not to be present, but they may emerge in future.

26. At the distribution level, small-scale storage, for example in batteries, can also improve quality of power provision.

ACKNOWLEDGMENTS

The authors would like to thank 3 anonymous reviewers and the special issue's editor Professor David Broadstock for insightful comments which have improved the clarity and quality of the paper. They acknowledge support from the EPSRC under grant EP/K002228. Monica Giulietti also acknowledges support from the EPSRC under grants EP/N001745/1 and EP/R062258/1 and from UKERC under grant UKERC/FF3/3. Elisa Trujillo-Baute also acknowledges support from the Generalitat de Catalunya (2015-SGR-531) and the Spanish Ministry of Economy (ECO2015-69107-R, MINECO/FEDER, UE). The authors would like to thank participants to the 5th meeting in International Economics at the University Jaume I in Villareal, Spain in July 2016 for useful comments on a previous version of this work, and several attendees at a seminar at the University of Barcelona where this work was also presented. Michael Waterson would like to thank the Chair of Energy Sustainability at the University of Barcelona for its generous hospitality during his visit in November 2015. Luigi Grossi wishes to acknowledge the Department of Economics at the University of Warwick for its hospitality in different periods during which most of the R scripts used in this paper were written.

REFERENCES

- Barton J. P. and M. Thomson (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.
- Barbry A., M.F. Anjos, and E. Delage (2017). "Robust self-scheduling of a price-maker energy storage facility in the New York electricity market." mimeo HEC Montreal, (<http://web.hec.ca/pages/erick.delage/BarbryAnjosDelage-May2017.pdf>)
- Bradbury, K., L. Pratson, and D. Patiño-Echeverri (2014). "Economic viability of energy storage systems based on price arbitrage potential in real-time U.S. electricity markets." *Applied Energy* 114: 512–519. <https://doi.org/10.1016/j.apenergy.2013.10.010>.
- Carson, R.T. and K. Novan (2013). "The private and social economics of bulk electricity storage." *Journal of Environmental Economics and Management* 66(3): 404–423. <https://doi.org/10.1016/j.jeem.2013.06.002>.
- Cruise, J., L. Flatley, R. Gibbens, and S. Zachary (2015). "Optimal control of storage incorporating market impact and with energy applications." arXiv 1406.3653v2.
- Cruise J., L. Flatley, and S. Zachary (2016). "Impact of storage competition on energy market." arXiv 1606.05361.
- ECOFYS (2014). "Energy Storage Opportunities and Challenges—A West Coast Perspective." White Paper [Online]. Available from: <http://www.ecofys.com/files/files/ecofys-2014-energy-storage-white-paper.pdf> [5 January 2015].
- Elexon Ltd, Market Index Definition Statement version 8, 23/2/2017 (www.elexon.co.uk/wp-content/uploads/2017/02/MIDS_v8.0.pdf)
- Esteban, M., Q. Zhang, and A. Utama (2012). "Estimation of the energy storage requirement of a future 100% renewable energy system in Japan." *Energy Policy* 47: 22–31. <https://doi.org/10.1016/j.enpol.2012.03.078>.
- Flatley, L., R.S. MacKay, and M. Waterson (2016). "Optimal strategies for operating energy storage in an arbitrage market." *Journal of Dynamics and Games* 3(4): 371–398. <https://doi.org/10.3934/jdg.2016020>.
- Giulietti, M., L. Grossi, and M. Waterson (2013). "The economic value of storage: evidence from the UK." University of Barcelona, Chair on energy and sustainability, Document de Treball, 2013/37, November.
- Gowrisankaran, G., S. Reynold, and M. Samano (2016). "Intermittency and the value of renewable energy." *Journal of Political Economy* 124(4): 1187–1234. <https://doi.org/10.1086/686733>.
- Grossi, L., S. Heim, and M. Waterson (2017). "The impact of the German response to the Fukushima earthquake" *Energy Economics* 66: 450–465. <https://doi.org/10.1016/j.eneco.2017.07.010>.
- Grossi, L. and F. Nan (2015). "Robust estimation of regime switching models." In *Advanced Statistical Models Data Analysis*, I. Morlini, M. Vichi, and T. Minerva (Eds.), chapter 14, 125–135, Berlin: Springer. https://doi.org/10.1007/978-3-319-17377-1_14.
- Hutchinson, A. (2015). Unpublished PhD Thesis, University of Warwick, December 2015.
- Lund, H., G. Salgi, B. Elmegaard, and A.N. Andersen (2009). "Optimal operation strategies of compressed air energy storage (CAES) on electricity spot markets with fluctuating prices." *Applied Thermal Engineering* 29: 799–806. <https://doi.org/10.1016/j.applthermaleng.2008.05.020>.

- Nasrolahpour, E., J.S. Kazempour, H. Zareipour, and W.D. Rosehart (2016). "Strategic Sizing of Energy Storage Facilities in Electricity Markets." *IEEE Transactions on Sustainable Energy* 7(4): 1462–72. <https://doi.org/10.1109/TSTE.2016.2555289>.
- Newbery, D. (2016). "Missing money and missing markets: reliability, capacity auctions and interconnection." *Energy Policy* 94: 401–410. <https://doi.org/10.1016/j.enpol.2015.10.028>.
- Renewable Energy Association (2015). "Energy Storage in the UK: An Overview." November [Online]. Available from: https://www.r-e-a.net/upload/rea_uk_energy_storage_report_november_2015_-_final.pdf [2 January 2018].
- Sioshansi, R., P. Denholm, T. Jenkin, and J. Weiss (2009). "Estimating the value of electricity storage in PJM: Arbitrage and some welfare effects." *Energy Economics* 31(2): 269–277. ISSN 0140-9883, <http://dx.doi.org/10.1016/j.eneco.2008.10.005>.
- Sioshansi, R. (2010). "Welfare Impacts of Electricity Storage and the Implications of Ownership Structure." *The Energy Journal* 31(2): 173–198. <https://doi.org/10.5547/ISSN0195-6574-EJ-Vol31-No2-7>.
- Steward, D., G. Saur, M. Penev, and T. Ramsden (2009). "Lifecycle Cost Analysis of Hydrogen Versus Other Technologies for Electrical Energy Storage, National Renewable Energy Laboratory." Technical Report NREL/TP-560-46719, November.
- Vatani, B., B. Chowdhury, S. Dehghan, and N. Amjadi (2018). "A critical review of robust self-scheduling for generation companies under electricity price uncertainty." *International Journal of Electrical Power & Energy Systems* 97: 428–439. <https://doi.org/10.1016/j.ijepes.2017.10.035>.
- Walawalkar, R., J. Apt, and R. Mancini (2007). "Economics of electric energy storage for energy arbitrage and regulation in New York." *Energy Policy* 35(4): 2558–2568. <https://doi.org/10.1016/j.enpol.2006.09.005>.
- Zafirakis, D., K.J. Chalvatzis, G. Baiocchi, and G. Daskalakis (2016). "The value of arbitrage for energy storage: Evidence from European electricity markets." *Applied Energy* 184: 971–986. <https://doi.org/10.1016/j.apenergy.2016.05.047>.