

FOR 1 2013

ISSN: 1500-4066

April 2013

Discussion paper

Now or Later? Trading wind power closer to real-time and how poorly designed subsidies lead to higher balancing costs

BY
Johannes Mauritzen

Now or Later?

Trading wind power closer to real-time and how poorly designed subsidies lead to higher balancing costs

Johannes Mauritzen
NHH Norwegian School of Economics
Department of Business and Management Science
Helleveien 30
5045 Bergen
Norway

and

Research Institute of Industrial Economics (IFN)
Stockholm, Sweden

johannes.mauritzen@nhh.no

JEL Codes: Q42; Q48

Abstract

An important challenge facing many deregulated electricity markets is dealing with the increasing penetration of intermittent generation. Simulation studies have pointed to the advantages of trading closer to real-time with large amounts of intermittent generation. Using Danish data, I show that, as expected, shortfalls increase the probability of trade on the short-term market. But in the period between 2010 and 2012 surpluses are shown to decrease the probability of trade. This unexpected result is likely explained by wind power policies that discourages trading on Elbas and leads to unnecessarily high balancing costs. I use a rolling-window regression to support this claim.

I have recieved financial support for this research from the Norwegian Center for Sustainable Energy Studies (CenSES) as well as the IFN research program The Economics of Electricity Markets. I would like to thank Henning Parbo, Signe Horn Rosted, Linda Rud, Thomas Tangerås, Pär Holmberg, Ewa Lazarczyk, Henrik Horn, Niels Vestergaard Jonas Andersson and Gunnar Eskeland for helpful comments and discussion.

1. Introduction

Many deregulated electricity markets, including the common Nordic market, have traditionally relied heavily on a day-ahead market mechanism where trade is organized between 12 and 36 hours ahead of actual delivery. The installation of large amounts of intermittent power sources such as wind power poses serious problems for this type of market mechanism.

The reason is of course that wind power cannot be scheduled and can only be forecast approximately. However this forecast becomes better the closer one gets to the time of delivery (Holttinen et al. 2006). Large amounts of wind power are therefore widely expected to lead to a heavier reliance on markets that trade closer to the time of delivery, like the Nordic hour-ahead market called Elbas (Nord Pool Spot, 2012).

A growing literature has developed around the issue of dealing with intermittency in deregulated power markets. Particularly relevant to this article is Holttinen (2005) who, using a simulation model, estimates that producers could reduce their balancing costs by 30% by trading on an hour-ahead market like Elbas as opposed to a day-ahead market. Holttinen et al. (2006) uses two commercial wind power forecasting tools and data from a Finnish wind power farm to calculate a balancing cost of between .6 and .8 EUR/MWh if power is bid in the market between 12 and 36 hours ahead of delivery. Yet to my knowledge, no empirical studies exist on the effects of intermittent energy and forecast error on the use of markets that trade closer to real time.

Previous empirical studies investigating the effects of wind power on balancing costs have used total amount of wind power as a proxy for the effects of forecast error (Morthorst (2003), Forbes and Zampelli (2007)). However I will show that the use of total amount of wind power can introduce a potentially severe bias in the estimation since it does not take into account the asymmetric effects of positive versus negative forecast errors. In this paper, I use data on day-ahead forecasted as well as settled wind power in western Denmark in order to create measures of positive and negative forecast errors.

Figure 1 shows the ex-ante expected relationship between the forecast error of wind power and the probability of trade on the Elbas market. Negative forecast errors - where actual wind power produced is less than that estimated a day ahead – can be expected to lead to a higher probability of trade on the Elbas market. As wind power producers realize that they will not be able to generate as much electricity as they had expected, they have a strong incentive to go on the Elbas market in order to make up for the shortfall. By doing so, they reduce the costs they incur in the balancing market.

The same logic should apply for positive forecast errors - where more wind power is delivered than forecast. Wind power producers have an incentive to go on the Elbas market to sell the excess electricity in order to avoid incurring balancing costs.

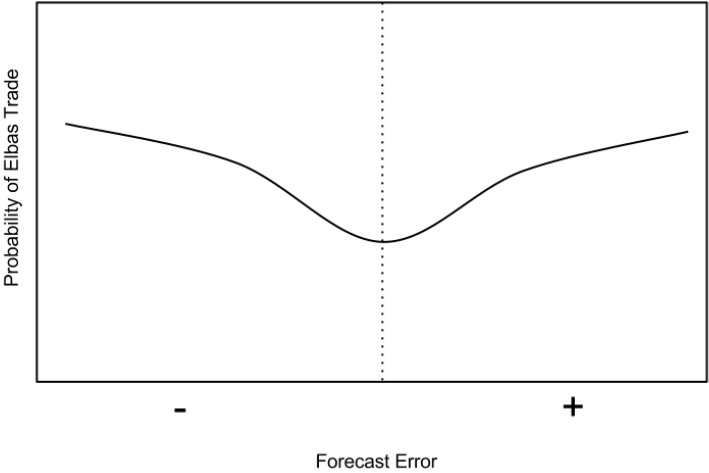


Figure 1: Ex-ante expected relationship between wind forecast error and probability of trade on the Elbas market.

Yet an empirical estimate of the relationship gives a more nuanced picture, as figure 2 shows. A shortfall of wind power production has the expected result of increasing the probability of trade on the Elbas market. This result provides strong support to the idea that trading closer to real-time can reduce the balancing costs associated with having large amounts of wind power and other intermittent generation. However, a surplus of wind power has the unexpected effect of reducing the probability of trade on the Elbas market.

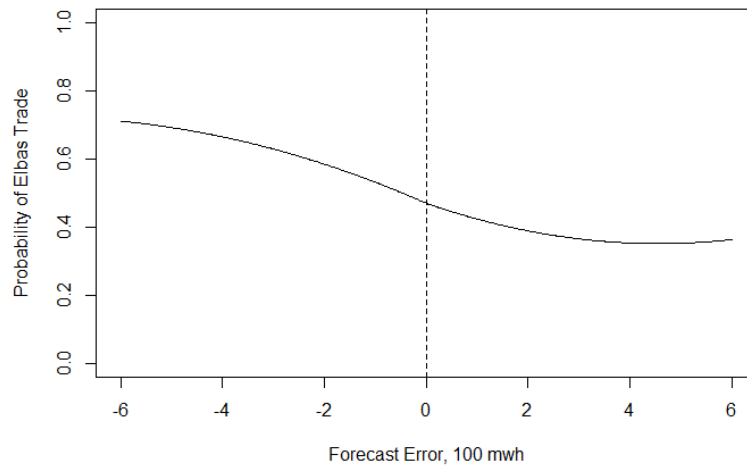


Figure 2: Estimated relationship between wind forecast error and probability of trade on the Elbas market

The most likely reason for this unexpected relationship is that subsidies and regulation intended to encourage investment in wind power have led to a perverse incentive for wind power producers to avoid the Elbas market when they produce surplus wind power.

I use a linear probability models to estimate both the effect of wind forecast errors on the probability of market trading. The models are specified to allow for a quadratic relationship between forecast errors and probability of market trade. They also allow for a non-linearity at zero – that is, the slopes are allowed to differ between positive and negative forecast errors.

It might be questionable whether such a simple linear model sufficiently captures the relationship between forecast errors and probability of trade. To give an idea of the appropriateness of the linear model, I also use a simple non-parametric estimation of the effect of forecast errors on the probability of market trade. This provides a visualization of the expected probability of trade given the forecast error. From this, the simple quadratic specification is shown to have a surprisingly good fit.

To give the results a causal interpretation I rely on the assumptions that the wind forecast error variables are exogenous and independent. Though such assumptions can often be strong, I argue that in this case they are justifiable. I also discuss potential violations of these assumptions and test for their relevance.

Finally as evidence for the claim that the unexpected negative relationship between positive forecast error and probability of Elbas trade can be explained by subsidy policy, I run a rolling windows regression with updated data through February of 2013. This shows that as the subsidies expired for a large amount of wind power over the course of 2012, the effect of positive forecast errors reversed and became positive, as originally expected.

2. Market, Data and Methodology

The Nordic market is a good testing ground for the effects of intermittency and forecast error on short-term market trading. The Nordic market is one of the oldest market-based electricity systems, dating back to the Norwegian electricity market reform of 1991, and is generally seen as being well functioning and efficient. For a thorough history and overview see Rud (2009). The Nordic market also has several market mechanisms: a day-ahead market – the so called “Spot” market, a continuous hour-ahead market called Elbas, and balancing and regulating markets operated at a national level.

Denmark, which became fully integrated in the Nordic market in 2003, has a relatively long history of feeding large amounts of wind power into its grid. In 2011 wind power made up approximately 27% of all electricity production in the country (Energistyrelsen, 2012). About 75% of this wind power is produced in the western Denmark price area, consisting of Jutland and Fyn.

The hour-ahead Elbas market began operation in Sweden and Finland in 1999. Eastern Denmark joined in 2004 and western Denmark joined in 2007. Elbas has later been extended to Norway, Estonia and northern Germany. In addition the market has been linked with the Dutch-Belgian intraday market.

The timing of trade in Denmark is shown in figure 3. Bids for the day-ahead “Spot”¹ market, operated by the Nordic central exchange, Nord Pool, must be received by noon the day before delivery. Producers and consumers submit bids for every hour of the following day and from these bids Nord Pool establishes virtual demand and supply curves. Prices for each hour are

¹ To avoid confusion, I will from now on refer to this market as the “day-ahead” market, though the official name is the Nord Pool Spot Market

determined by the intersection of these curves. In Denmark two price areas exist – Denmark East and Denmark West. Prices in these areas diverge from the Nordic wide system price if congestion occurs on the transmission net.

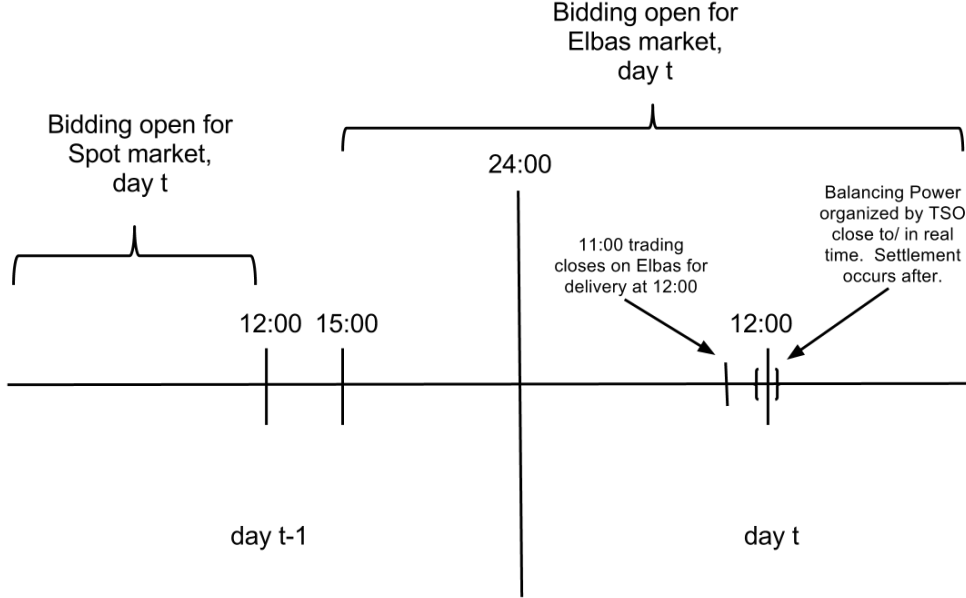


Figure 3: The time series of the wind power forecast error is centered around zero and can be shown to be stationary.

In Denmark, Elbas is opened for trade at 15:00 the day before delivery and continues up to an hour before delivery. In the figure, 12:00 in day *t* is shown as the time of delivery purely for illustrative purposes. In practice trading on the Elbas market can happen between an hour to at most 32 hours ahead of delivery. Unlike the day-ahead market, no single settled price exists for each hour since the Elbas market operates continually – buyers and sellers directly choose to accept bids placed on the market.

Finally a balancing market exists in order to deal with any imbalances close to real time. The Danish balancing market is operated locally by the Danish transmission system operator, Energinet, but is integrated with the Nordic wide system called NOIS. Producers make bids a day ahead² to provide either upward regulation – where they add power to the grid or downward regulation – where they reduce power to the grid. Energinet then activates the bids in order of price. The price for any given hour is established by the highest activated bid. The day-ahead spot price is regulated to be the minimum regulation-up price and the

² Producers can also adjust their bids up to 45 minutes before the time of delivery

maximum regulation-down price. For more details on the balancing market see Energinet (2008)

Unlike the day-ahead market, activity on the Elbas market is relatively light and not enough to support trading at every hour. In the sample of data I use in the years 2010 and 2011, trading happened only in 46% percent of the hours. Given the supposed advantages of trading closer to real-time, this can seem a bit puzzling. Weber (2010) takes up this question by looking at issues of liquidity in intraday markets. The author argues that without sufficient liquidity, a potential bidder may be able to influence the price in the market which in turn can act as a major transaction cost for the participant. The author also notes that a continuous trading system – such as the Elbas market – does not have the same price transparency as day-ahead markets. In a similar vein, day-ahead markets are able to aggregate liquidity into hourly auctions rather than individual trades.

On the other hand, Energinet explains the relative lightness of trade on Elbas compared to the day-ahead market as being due to the higher transaction costs involved. Trading on the Elbas market requires activity and market surveillance over 24 hours. Given how close the prices generally are in the day-ahead, Elbas, and balancing markets, it generally makes more sense to bid into the day-ahead market and let the transmission system operator deal with imbalances by way of the balancing market. At the same time, market actors highly value being able to use the Elbas market to correct for mistakes or unexpected events (Rosted 2012).

I use hourly data on both wind forecast error and trading on the Elbas market in the years 2010 and 2011 – a total of 17,520 observations. The data on estimated and settled wind power as well as on trading and prices on the Elbas market is publicly available on the website of Nord Pool Spot (www.nordpoolspot.com) and the Danish transmission system operator, Energinet (www.energinet.dk). A cleaned dataset as well as the R code for the full analysis is available on my website at:

<https://sites.google.com/site/johannesmauritzen/home/publications>

I only use data as far back as the beginning of 2010 because Nord Pool only started publishing data on both realized and estimated wind power production in Denmark in late 2009. Previously only realized wind power production was published. It would have been

interesting to investigate the effect that wind power forecast errors have on turnover as well. Unfortunately, turnover data for the Elbas market is not published by Energinet or Nord Pool.

The data on estimated wind power production is from Energinet's forecasting model that they run and publish the day before delivery. The five large power companies that are responsible for trading their own wind power as well as that of smaller producers have their own forecasting models that they use to bid into the day-ahead market. However, neither their bids nor the results from their models are publicly available. Yet the estimate from Energinet will still likely serve as a good unbiased estimator of day-ahead expected wind power. Energinet could easily detect a persistent bias between its own estimates of wind power production and that bid into the day-ahead market. They could then either investigate any irregularities or adjust their model.

The time series of wind power forecast errors is plotted in figure 4. The series is centered around zero – or to be more precise a mean of zero cannot be rejected. The series can also be shown to be stationary by way of an Augmented Dicker-Fuller (ADF) test. Stationarity – in a very simplified sense the idea that the mean and variance of a series is constant - is important in establishing the validity of the coefficients estimated below. For further details see, for example, Hamilton (1994).

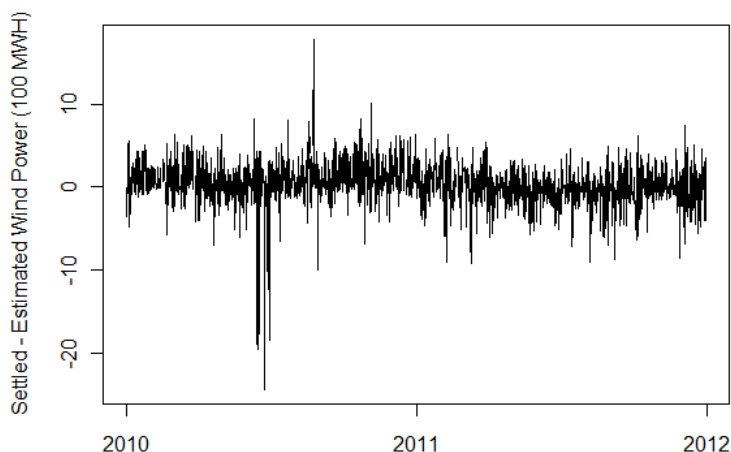


Figure 4: The time series of the wind power forecast error is centered around zero and can be shown to be stationary.

The empirical probability model I use is simple, but flexible enough to allow for two likely non-linearities. The first is an inflection point at zero for the effect of the forecast error. The relationship between a positive error and a negative error are unlikely to be completely linear, and I allow for different slopes in the positive and negative direction. Since the left-hand side is a probability, we would also expect the effects of the forecast error to curtail at a certain point. At the extreme, it would of course be nonsensical to have a probability of above 100%. Therefore I also include quadratic terms for both the effects of wind error in the positive and negative error.

The simple model can be written as in equation 1.

$$Prob_t^{Elbas} = \alpha + \beta_1 forError_t^+ + \beta_2 (forError_t^+)^2 + \beta_3 forError_t^- + \beta_4 (forError_t^-)^2 + \epsilon_t \quad (1)$$

Here $Prob_t^{Elbas}$ represents the probability of trade on the hour-ahead Elbas market for each hour, t . $forError_t$ represents either the positive (+) or negative (-) forecast error. These are of course mutually exclusive. ϵ_t represents the error term and α and the β 's represent parameters to be estimated.

Because I estimate both the positive and negative forecast error in a single equation I am implicitly assuming constant variance. However, estimating the effects of positive and negative forecast errors separately, and thus allowing for differing variances, does not significantly change the results.

Because of the properties of the wind forecast error variable that I use, I claim that I can give my results a causal interpretation. My identifying assumption is that the wind forecast error is both exogenous to trading activity on the Elbas market and that the forecast error is independent of other factors that could affect the probability of trading activity.

In many contexts exogeneity and independence of a variable can be strong assumptions. However in this case, they likely hold. The wind forecast error is the difference between day-ahead expected wind power production and actual realized wind power production. In the short-term, with a relatively stable amount of installed wind power capacity, both of these are fundamentally a function of the wind speed and location. The speed and location of wind is of course not something that can react to market conditions or prices. More so the marginal

production costs of wind power producers are very low– they have little reason to hold back power even at low prices.

A few possible exceptions to the exogeneity of wind power are worth mentioning. Though actual wind production may be exogenous, wind power producers may have an incentive to not bid their expected value of wind power on the market (see Rud (2009)). Systematic under or overbidding, though, is likely easily identifiable by Nord Pool or the transmission system operator.

The Nordic market is characterized by dominant producers at the national level, and in Denmark, Dong Energy has an approximately 50% share of the power market. Presumably such a company with a range of generation technology might have the incentive to hold back wind power in order to generally increase prices. Again, this would likely be relatively easy for the transmission system operator to detect. More so, given the relatively generous feed-in tariffs that wind power receives in Denmark, it is difficult to think of a scenario where producers would willingly withhold wind power production.

The assumption that wind power forecast errors are independent can be argued in a similar vein as its exogeneity. Even if a factor was correlated to actual wind power production – such as other weather conditions like rain or cold that could affect conditions on the market – it is unlikely that they would be correlated with the forecast error. The forecast error can, in other words, be considered a truly random event that can be seen as an exogenous “shock” on the operation of electricity markets.

A few plausible exceptions exist for the independence assumption. The first is total amount of wind power, which is shown to be correlated with wind forecast errors. If the amount of wind power has an effect on Elbas trading on top of the effect of increased forecast errors, then the exclusion of this variable could lead to biased estimation. Another potential factor that could both affect trading on the Elbas market as well as be correlated with wind forecast errors could be consumption forecast errors. For example, unexpectedly windy weather might lead to both a surplus of wind power as well as more consumption than was expected as more people stay inside and use appliances. Luckily, data exists for both total amount of wind power and consumption forecast errors, and the inclusion of these factors cannot be shown to affect the estimation.

It is also worth briefly discussing the meaning of the error term, ϵ , in these models, and in turn the idea of a *probability* of trading on the Elbas market. The idea is not to imply that trading on the Elbas market is an inherently random process, like rolling a dice. Instead the error term contains a host of both observed and unobserved factors that affect whether trade happens or not. But as discussed these factors are argued to be independent of the forecast error and therefore their omission as explicit factors in the model will not bias the results.

Factors such as expectations of prices on the balancing market and day-ahead market as well as congestion on the transmission grid all likely play a role in determining the amount of trade on Elbas and the interaction of these variables is an interesting and important research question. Though some discussion of these factors is provided, a thorough investigation of the interaction of such factors likely requires at a minimum structural econometric modeling and/or simulation type modeling, and is well outside the narrow scope of this article.

To give an idea of the conditional relationship between variables as well as to check the appropriateness of my linear model, I use a simple non-parametric technique. Instead of estimating probabilities by way of a parametric model, I take the means of ordered sub-samples – or “bins” - in effect creating a series of conditional expectation estimates of the dependent variable at different values of the independent variable. I then use a kernel smoother - or non-parametric regression - through the estimated conditional expectation estimates. My intent is not to provide a rigorous non-parametric regression, but instead to simply help motivate the linear model and provide a visualization of the relationship between variables. Therefore I leave out much of the technicalities of non-parametric regressions such as bandwidth calibration, curve-wise standard error estimation, etc.

Logit or probit models are popular choices when using binary dependent variable data. The estimates from these models can be somewhat more efficient compared to those from linear probability models if the implied structural assumptions imposed are true. Yet those functional and distributional assumptions can often be difficult to justify in practice. I choose to use the simpler linear probability model, which when estimated by ordinary least squares (OLS) has a straightforward interpretation and is always the minimum mean squared error linear estimate of the underlying conditional expectations function (Angrist and Pischke 2009).

3. Results

Table 1 shows the results from the regressions of the probability of trade on the Elbas market. The first column of the table shows the results from where I naively try to estimate the effect of total amount of wind power generated in western Denmark on the probability of trade on the Elbas market. The coefficient on total amount of wind power, labeled *TotalWind*, cannot be shown to be significantly different from zero. This might suggest that wind power has no effect on the probability of trade in the Elbas market. In turn one might come to the conclusion that trading closer to real-time does not have a significant effect on reducing balancing costs. However this regression suffers from a severe bias.

Table 1. Probability of Elbas Trade and Wind Power Forecast Errors

	1	2	3	4
<i>Intercept</i>	0.466 (0.000)	0.475 (0.000)	0.469 (0.000)	0.468 (0.000)
<i>TotalWind</i>	0.000 (0.816)	n/a	0.002 (0.370)	0.002 (0.405)
<i>error_neg</i>	n/a	0.066 (0.000)	0.064 (0.000)	0.065 (0.000)
<i>error_neg_sq</i>	n/a	-0.004 (0.000)	-0.004 (0.000)	-0.004 (0.000)
<i>error_plus</i>	n/a	-0.052 (0.000)	-0.056 (0.000)	-0.055 (0.000)
<i>error_plus_sq</i>	n/a	0.006 (0.003)	0.006 (0.003)	0.006 (0.003)
<i>cons.error.plus</i>	n/a	n/a	n/a	-0.001 (0.899)
<i>cons.error.neg</i>	n/a	n/a	n/a	0.001 (0.701)

17436 observations

P-values are in parenthesis, based on Newey-West standard errors

Right hand variables are in 100 mWh units

In principle, large amounts of wind power can be handled by a day-ahead market mechanism with little problem – if it is accurately forecasted. In practice, the amount of wind power and the size of the forecast errors are correlated. Figure 5 shows how wind power is positively

correlated with both negative and positive forecast errors. Thus the use of total amount of wind power as a proxy for forecast error implicitly acts as if the sign of the forecast error is not important. In turn, using total amount of wind power can only serve as a reliable proxy for the effects of forecast error if the effect on the variable of interest is symmetric. This is certainly not the case for activity on the Danish Elbas market, as will be shown.

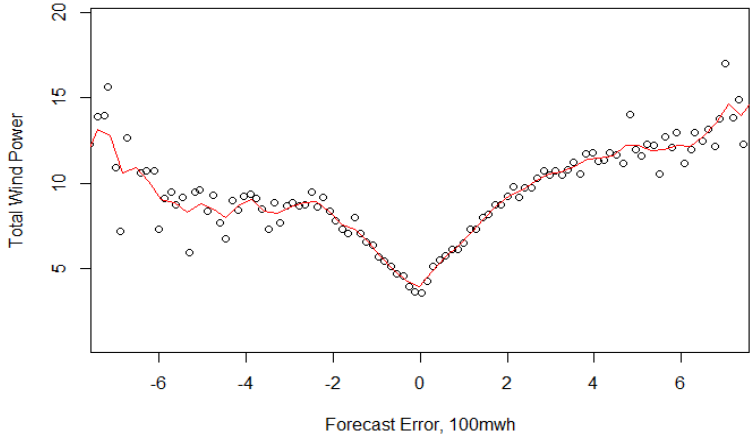


Figure 5: The relationship between wind power forecast errors and total amount of wind power.

The second column of the table shows the results from the regression with negative and positive forecast errors and their square terms. The fitted model can then be written as in equation 2.

$$Prob_t^{Elbas} = .47 - .052forError_t^+ + .0057(forError^+)_t^2 + .066forError_t^- - .0043(forError^-)_t^2 + \epsilon_t \quad (2)$$

The smooth black line in figure 6 shows the fitted model. As discussed earlier, the results indicate that a negative forecast error leads to a higher probability of trade on the Elbas market but a positive forecast error has the unexpected effect of reducing the trade on the market. The magnitudes of these effects are quite similar, which explains why a regression on total amount of wind power led to an estimated coefficient of zero. The effects of negative and positive forecast error, both positively correlated with total amount of wind power, cancel each other out in this regression.

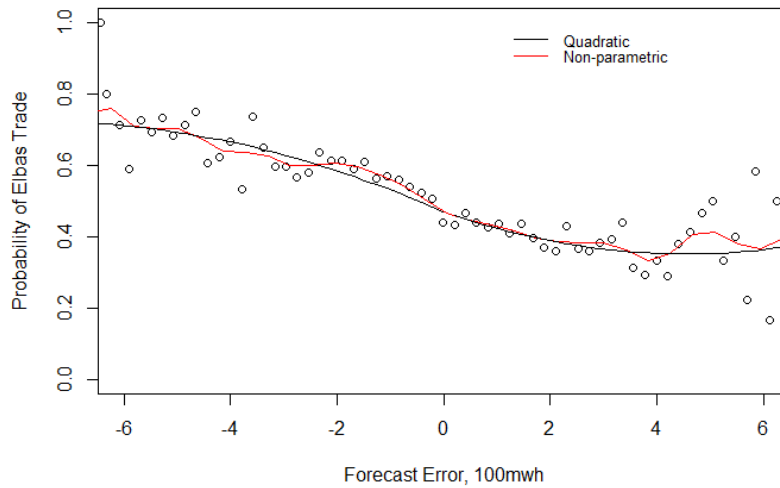


Figure 6: The kernel smoothed nonparametric estimation of the effects of forecast error indicates that the quadratic model has a surprisingly good fit to the data.

Clearly, we cannot extrapolate from the estimated quadratic function too far in either direction. Taken literally the function implies that beyond a certain negative forecast error the probability of trading on the Elbas market would actually decrease and vice versa for positive forecast errors. With the above estimated coefficients this would occur just below -600 mWh and just above 700 mWh. Still, the vast majority of forecast errors in the sample happened between -600 mWh and 600 mWh as the histogram in figure 7 shows. In this respect, the quadratic form seems adequate.

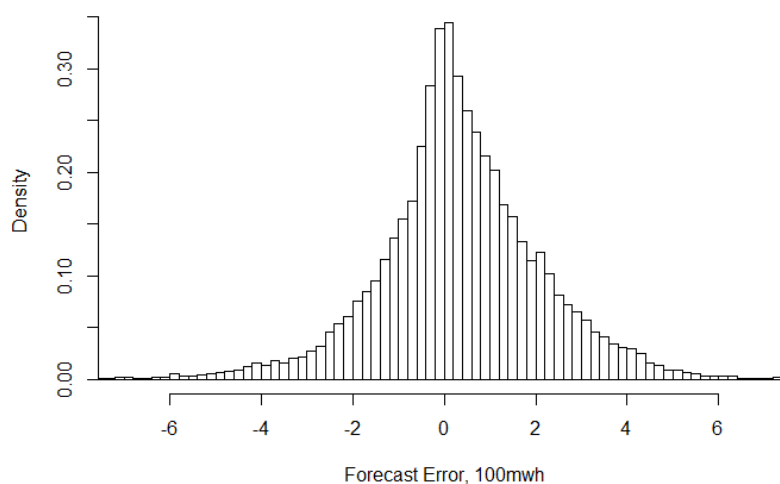


Figure 7: The vast majority of forecast errors are between -600 and 600 mWh. The quadratic model is therefore likely adequate as it provides sensible probability estimates within this range.

The estimated coefficients of the simple quadratic model are all highly statistically significant. Yet some doubt is warranted on how well-fitting this simple model is in describing the relationship between wind forecast errors and trading on the Elbas market. It is not hard to imagine more complex or more non-linear relationships. For example, a threshold level of forecast error that leads to a jump in Elbas market trading.

Because of the simplicity of the quadratic specification, the appropriateness of the model is easily verified by using the simple non-parametric technique described in the previous section. Referring back to figure 6, the small circles are the conditional expectation estimates of the probability of Elbas trade from each forecast error bin. For example, one point might represent the average of all the observations where the wind power forecast error was between -100 and -105 mWh. The red line is the kernel smoothing function through these points. These are both overlaid on the fitted parametric model estimated above, which is the black line. By inspection, the fit of the quadratic model appears surprisingly good.

The dispersion of the conditional expectation estimates visible towards the ends of the function come from the fact that there are a much higher number of observations per bin closer to zero. The points towards the ends are then estimated with higher variance.

The results do not carry over to eastern Denmark (see table 2 in the appendix) where no significant relationship between forecast error and probability of trade in the eastern Denmark price area can be estimated. Western Denmark contains approximately 75% of Denmark's wind power capacity as of 2012, thus wind power makes up a significantly larger part of total generation capacity in the western part of Denmark compared to the eastern Denmark.

I argue that the reason I can use such a simple model to estimate the effects of wind forecast errors on Elbas trade is that the forecast errors are unlikely to be correlated with any factors that in turn are also correlated with trade on the Elbas market. The wind forecast errors can be seen as random exogenous shocks independent of other factors. Two plausible exceptions to this may exist and could potentially introduce a missing variable bias. The first, as has been discussed, is total amount of wind power, which is positively correlated with the forecast errors. In the third column in the table I include both total amount of wind power as well as

the forecast errors. The inclusion of total amount of wind power has no substantial effect on the estimated coefficients on the forecast error terms. Furthermore, the estimated coefficient on wind power itself is not significantly different from zero. It is worth repeating: wind power generation in itself can be handled on a day-ahead market, it is the associated forecast errors that can lead to higher balancing costs and the desirability of trading closer to real time.

The other potential exception to the independence of the forecast error variables is if they are correlated with consumption forecast errors, which in turn may have an effect on trading in the Elbas market. In the fourth column I therefore include a measure of positive and negative consumption errors for Denmark. These are created using Energinet's day-ahead forecast of consumption as well as settled consumption. The regression results indicate that these variables do not affect the estimated coefficients on the wind forecast error variables. Furthermore, the estimated coefficients on the consumption error variables are not significantly different from zero.

4. Explaining the asymmetry: internal displacement, the balancing market and subsidy policy

To a certain extent, an asymmetry in the effects that positive and negative wind power forecasting errors has on trading in the Elbas market was to be expected. As the decision tree in figure 7 shows, whether or not to go on the Elbas market is potentially complex, involving factors such as a power company's portfolio of power plants and prices on the day-ahead and balancing markets.

One source of potential asymmetry between the effects of positive and negative forecast errors is the ability of large power companies to displace their own production. Much of the wind power in Denmark is owned by large power companies with a diversified portfolio of generation. More so, small wind power producers generally do not trade on the markets directly, but instead contract out trading to one of the large "balancing responsible" power companies. These power companies then have an ability to internally balance any surpluses or deficits of wind power. However it may be easier to internally deal with a surplus – by turning off a gas-powered unit for example, than dealing with an unexpected deficit of power.

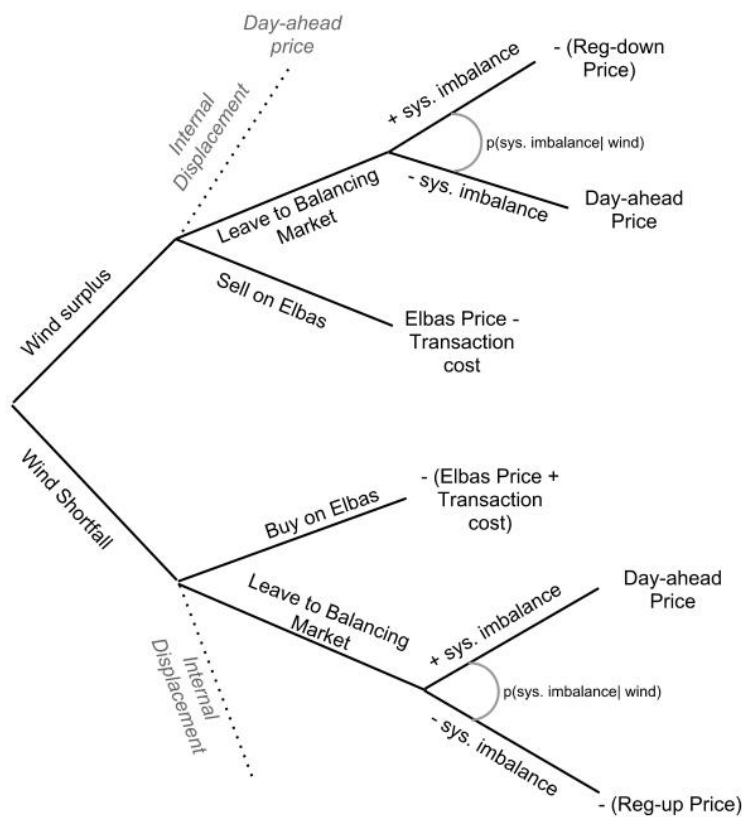


Figure 8: The decision of whether to go on the Elbas market can be complex, involving both prices in the day-ahead and balancing market, as well as power companies' ability to deal with imbalances internally.

Another source of potential asymmetry is the relationship between the Elbas market and the balancing market. The price on the balancing market – which the producers do not know ahead of time – determines the penalty that the producers pay for their imbalance. The balancing market is organized as a two-price model. If a producer's imbalance is in the same direction as the total imbalance, then they must pay the balancing market price – either the up-regulation price or the down-regulation price. However, if the producer's imbalance goes in the opposite direction of the system imbalance, in effect helping to resolve the system imbalance, they will receive the day-ahead price. Thus the penalty that a producer expects to pay for an imbalance is affected by both their expectations of what the price will be on the balancing market as well as the probability that that their imbalance is in the same or opposite

direction as the system imbalance. Complicating the picture is the fact that prices on the balancing market are themselves dependent on forecast errors, as figure 8 shows.

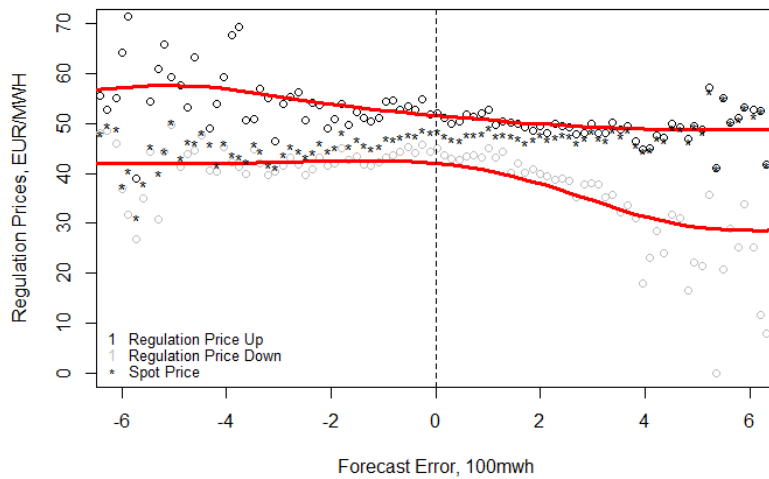


Figure 9: Regulation prices also are affected by wind forecast errors.

A full analysis of the dynamics of the Elbas market with the day-ahead and balancing markets would require a structural econometric or simulation approach and is well outside the scope of this paper. However, the discussion should make clear that though these factors may lead to an asymmetry in the magnitude of the effect of positive versus negative forecast errors, it does not explain why positive forecast errors are associated with a *lower* probability of trade on the Elbas market.

But a simpler explanation exists: the Electricity Supply Act of 1999 covers all turbines built before 2003, when new rules came into place. The act obligates the transmission system operator to purchase all wind power production for a period of 10 years from the date the turbine was connected to the grid (Energinet 2007). Importantly, the producers do not need to pay for the balancing costs associated with their production. Even after the 10 year period has expired and wind power producers must trade on market terms, the transmission system operator is still obligated to deal with balancing costs associated with production.

The reason behind the asymmetry between the effects of positive and negative forecast errors on Elbas trade now becomes clear. With a shortfall of wind power, a producer is still responsible for the balancing costs – the purchase obligation sensibly does not cover expected wind power. Thus a producer with a shortfall of wind power has an incentive to go on the

Elbas market in order to avoid the penalty on the balancing market. However with a surplus of wind energy, producers have a guaranteed buyer and face no balancing costs. They have no incentive to go on the Elbas market and the overall liquidity on the Elbas market decreases.

Luckily, this explanation provides a testable implication. Subsidies for turbines were lowered in 2003, leading to a large jump in installations in late 2001 and especially 2002. Turbines built in this period benefited from Energinet's purchase obligation for wind power, but only for 10 years. The sharp increase in installed wind power that could take advantage of the purchase obligation approaching 2003 in turns means a sharp decrease approaching 2013. We would then expect that the relationship between positive forecasting error and the probability of trade on the Elbas market would change as we near 2013 as gradually more of the wind turbines built over the course of 2002 lose the subsidy. In particular, we would expect, at a minimum, for the effect of positive forecast errors to approach zero if not move positive.

To test this prediction I extend the data on trade on the Elbas market as well as forecasting error through February of 2013. I then run a rolling-window regression of probability of trade on the Elbas market on positive and negative forecast errors. To simplify, I drop the quadratic terms, though including these terms does not materially affect the conclusions. Each window is composed of one year of data, or equivalently 8760 hourly observations. For each iteration, I move the window forward by 24 hours. This means that I run a total of 790 separate regressions.

The estimated coefficient on the positive forecast error for the window regressions is plotted along with the midpoint of the regression window in Figure 10. Clearly the relationship between positive forecast error and probability of Elbas trade begins to shift upwards as the regression window begins to include data from 2012. As the midpoint of the regression window moves into 2012, the estimated coefficients on positive forecast errors are estimated to be significantly positive and of a magnitude between .02 and .04. This can be interpreted to mean that for every 100 MWh increase in positive forecast error there is between a 2 and 4% increase probability of trade on the Elbas market.

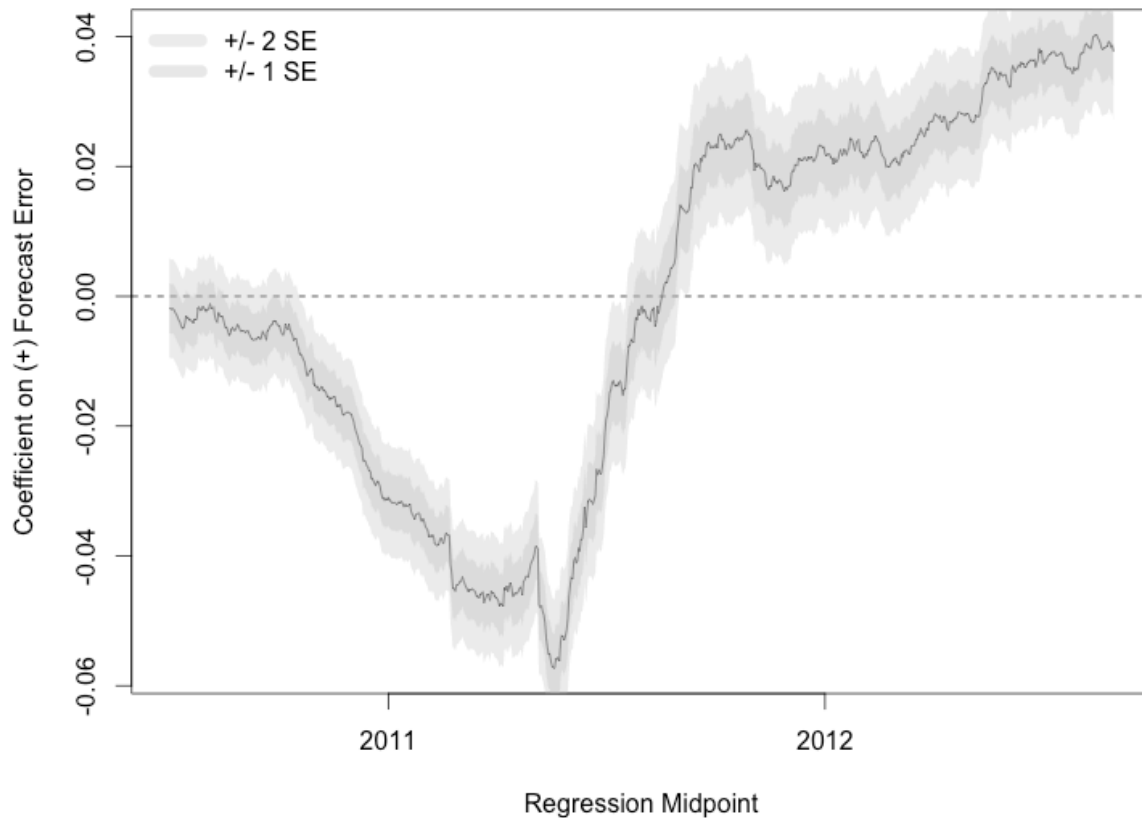


Figure 10: The coefficient on the positive forecast error term in a rolling window regression with window length of 1-year, or 8760 hourly observations. The shaded regions represent +/- one and two standard errors, which can be interpreted as approximately 70% and 95% point-wise confidence intervals.

Discussion and Conclusion

Both the arguments and methodology of this article are relatively simple, yet three important results emerge. The first is methodological. Studies of the effects of intermittency on power markets that have used total amount of wind power as a proxy for forecast errors are likely introducing a bias in their estimation. The reason is because total amount of wind power is positively correlated with both negative and positive forecast errors. However the effect that negative and positive forecast errors have on market operation is likely to be asymmetric. The results from the Elbas market are likely an extreme example of this, but the effects that surpluses and deficits will have on balancing markets and other areas of market operation are also likely asymmetric.

The second result is to show that a deficit of wind power compared to what was forecasted a day ahead will increase the probability of trade on the Elbas market. This provides strong

empirical evidence that the option of trading closer to the time of delivery can mitigate balancing costs associated with increased amounts of wind power.

The final result is to show how a poorly designed subsidy scheme for wind power has inadvertently lead to unnecessarily high balancing costs by eliminating the incentive to use the short-term market when surplus wind power is produced. While this result may seem to be limited to the particular design of the Nordic electricity market and Danish subsidy schemes, the implications are broader. Such purchase obligation incentives for renewable energy are a widely used investment incentive³. A country with a deregulated electricity market wishing to increase investment in wind power will likely be better off simply increasing the direct production subsidy, while keeping in place the market incentives for producers to deal with the balancing costs associated with forecasting errors.

Works Cited

Anon. "Nord Pool Spot." <http://www.nordpoolspot.com/How-does-it-work/Intraday-market-Elbas/>.

Energinet. 2007. *Regulation E : Settlement of Environmentally-friendly Electricity Generation*. <http://energinet.dk/SiteCollectionDocuments/Engelske dokumenter/EI/Regulation E Settlement of environmentally friendly electricity generation.pdf>.

———. 2008. *Regulation C2 : The Balancing Market and Balance Settlement*. <http://energinet.dk/SiteCollectionDocuments/Engelske dokumenter/EI/Regulation C2 The balancing market and balance settlement.pdf>.

Forbes, Kevin F, and Ernest M Zampelli. 2007. "Wind Power and the Electricity Market for Upward Regulation : Evidence for Western Denmark": 1–11.

Holtttinen, H. 2005. "Optimal Electricity Market for Wind Power." *Energy Policy* 33 (16) (November): 2052–2063. doi:10.1016/j.enpol.2004.04.001. <http://linkinghub.elsevier.com/retrieve/pii/S0301421504001016>.

Holtttinen, Hannele, Pirkko Saarikivi, Sami Repo, Jussi Ikäheimo, and Göran Koreneff. 2006. "Prediction Errors and Balancing Costs for Wind Power Production in Finland." In , 1–11. http://www.ieawind.org/AnnexXXIVSecure/MeetingInfo/Tasmania_2006/Holtttinen Windpred_Adelaide2006.pdf.

³ Purchase obligations can also refer to laws requiring that a certain percentage of a country or state's electricity production comes from renewable energy. To be clear, here I am referring only to the obligation of a transmission system operator to purchase all electricity produced by existing wind power.

- iMind. 2012. “Månedstatistik”. Energistyrelsen. http://www.ens.dk/dk/Info/TalOgKort/Statistik_og_noegletal/Maanedstatistik/Sider/Forside.aspx.
- Morthorst, P. E. 2003. “Wind Power and the Conditions at a Liberalized Power Market.” *Wind Energy* 6 (July): 297–308.
- Rosted, Signe Horn. 2012. “Correspondence My Email.”
- Rud, Linda. 2009. *Essays on Electricity Markets. Research in Economics*. http://www.snf.no/Files/Filer/Publications/Rapport/09/R10_09/R10_09.pdf.
- Weber, Christoph. 2010. “Adequate Intraday Market Design to Enable the Integration of Wind Energy into the European Power Systems.” *Energy Policy* 38 (7) (July): 3155–3163. doi:10.1016/j.enpol.2009.07.040. <http://linkinghub.elsevier.com/retrieve/pii/S0301421509005564>.

Appendix

	W. Denmark		E. Denmark
error_plus	-.052 (.012)	-.056 (.012)	.012 (.015)
error_plus ^2	.0058 (.0019)	.0060 (.0020)	-.0009 (.0009)
error_neg	.066 (.010)	.063 (.010)	-.040 (.045)
error_neg ^2	-.0042 (.0009)	-.0041 (.0008)	.012 (.023)
Tot. Wind	n/a	.0016 (.0018)	n/a
Intercept	.47 (.013)	.47 (.017)	.50 (.012)

Newey-West standard errors in parenthesis. 17236 Observations

Table 2. Effect of Wind Forecast Error on Probability of Elbas. Significant results can only be estimated for the western Danish price area. Forecast Error is in 100 mw units.