Whitepaper


**Hjalmar Pihl**

**November 2019**

hjalmar.pihl@ri.se | energyeconomicslab.se


# Swedish FCR prices – an analysis of the data

The purpose of this study is to gain a better understanding of when and why the prices for FCR services in Sweden increase. This is done by analyzing to what extent FCR prices correlate with – and can be predicted by – a range of electricity market variables.

## What is FCR?

The amount of electricity consumed and produced must, in every moment, be equal. To ensure that this balance is maintained, the transmission system operator (TSO) buys balancing services from producers and consumers that can quickly increase or decrease their production or consumption. The fastest-responding type of balancing service currently used in Sweden is called FCR (Frequency Containment Reserve). Resources that deliver this service respond directly to deviations from the 50 Hz target grid frequency. Such frequency deviations occur whenever there is a mismatch between production and consumption.

In Sweden, the TSO Svenska Kraftnät procures two types of FCR services: FCR-N (Normal) and FCR-D (Disturbance). FCR-N is used to continuously respond to both positive and negative deviations from 50 Hz, whereas FCR-D only responds to large negative frequency deviations. A producer that provides FCR-N therefore must be able to both increase and decrease its output, whereas a producer that provides FCR-D only needs to be able to increase production.

Svenska Kraftnät procures these services separately for each hour through competitive auctions that take place 1-2 days in advance. The resources that are willing to supply the services at the lowest price are awarded the contracts and the compensation is pay-as-bid. In practice, almost all FCR services in Sweden are currently supplied by hydroelectric power stations.

## Description of FCR prices

Because FCR is procured using pay-as-bid auctions, there is no single market clearing price for these services. Instead, Svenska Kraftnät publishes the weighted average amount paid to resources that supplied the services for each hour. For ease of exposition, we will refer to these amounts as FCR prices. However, it is important to keep in mind that they are not uniform market clearing prices.

In the following analysis, we use hourly data from January 2016 through April 2019, yielding almost 30 000 hourly observations. During this period, the average price for FCR-N was about 29 €/MW while the average price for FCR-D was lower – about 11 €/MW. Although both of these services are quoted in €/MW, they represent somewhat different quantities: a resource offering 1 MW of FCR-N capacity must be able to both increase and decrease its production (or consumption) by 1 MW, while 1 MW of FCR-D capacity only requires the resource to be able to increase its production (or decrease its consumption) by 1 MW.

The FCR prices have varied considerably around their means. The price distributions for both services are positively skewed with median prices below the mean, and occasional spells with prices far exceeding the means. Figures 1 and 2 below show how the FCR prices varied over the sample period, for FCR-N and FCR-D respectively. Each observation is here represented by a small rectangle, with dates along the x-axis and hours along the y-axis. The price is indicated by the color of the rectangle.

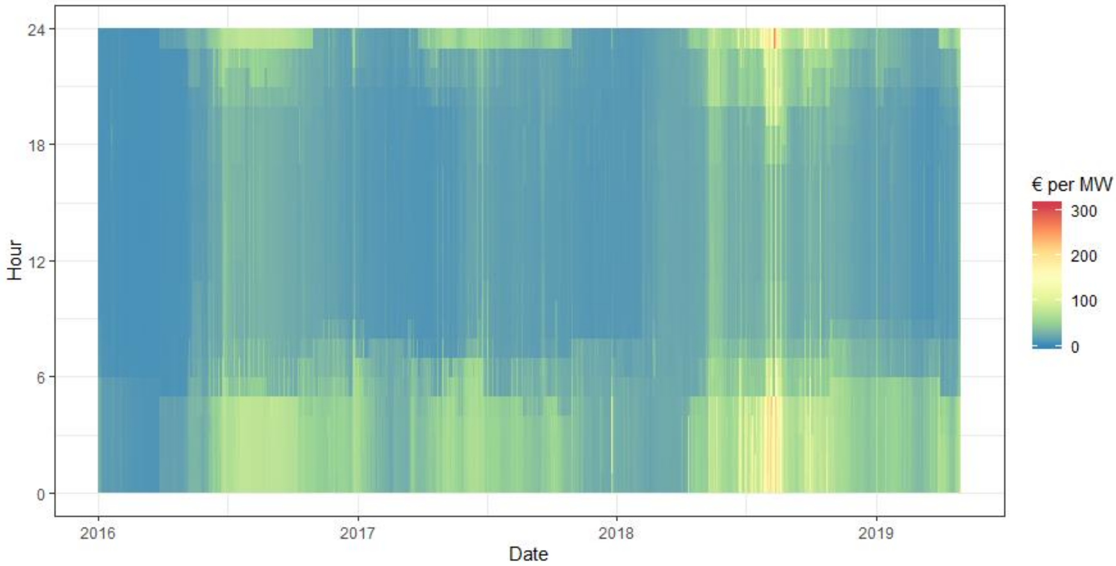


**Figure 1. FCR-N prices**

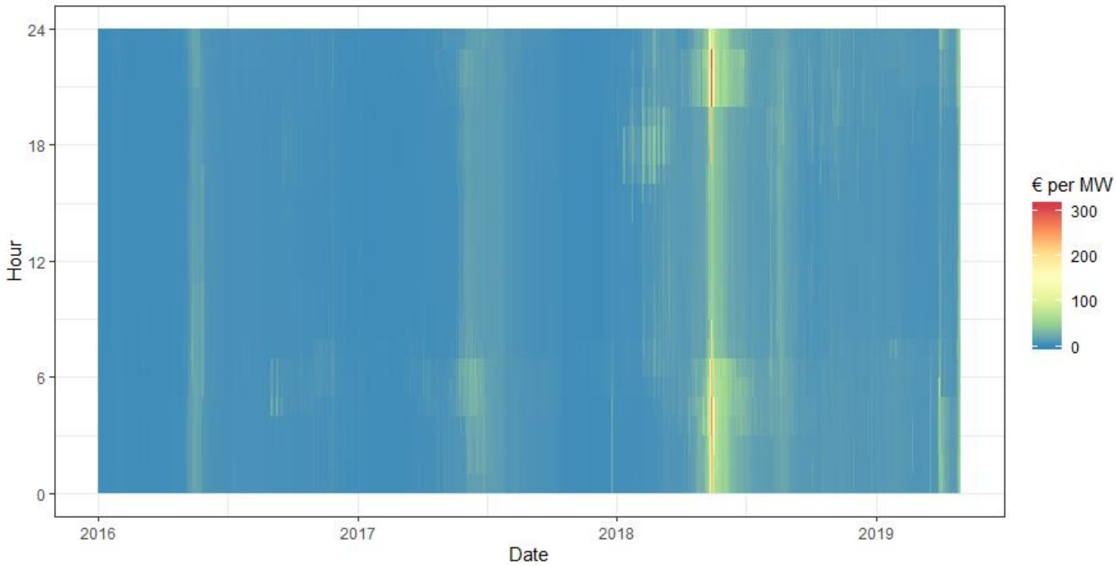


**Figure 2. FCR-D prices**

Figure 1 shows how FCR-N prices tend to be higher at night, and higher in the summer than in the winter. FCR-N prices are also typically somewhat higher during weekends, although this is more difficult to see in the figure. Prices for FCR-N therefore behave very differently compared to wholesale electricity prices, which are usually higher when consumption is high. Figure 1 also shows how the summer of 2018 experienced much higher prices than 2016 and 2017.

Figure 2 shows the same information for FCR-D. As seen, FCR-D prices remain relatively low for most of the time, with occasional periods of higher prices. There is less of a clear difference between night and day, and between weekdays and weekends. The periods of higher prices have mostly occurred in the late spring and early summer months. Again, 2018 stands out as having experienced exceptionally high prices.

# Factors that influence the price

We now turn to analyzing some factors that may help explain the price patterns shown above.

## Production from hydro

Because hydroelectric generators account for most FCR provision, factors affecting these resources are likely to influence the FCR prices. Figure 3 shows the association between the total amount of electricity generated by hydroelectric resources in Sweden for each hour and the FCR price of the same hour (FCR-N on the left and FCR-D on the right). This clearly shows that high FCR-N prices tend to occur when production from hydro is low, but that FCR-D prices do not appear to have an as clear of a connection to the amount of hydro production.

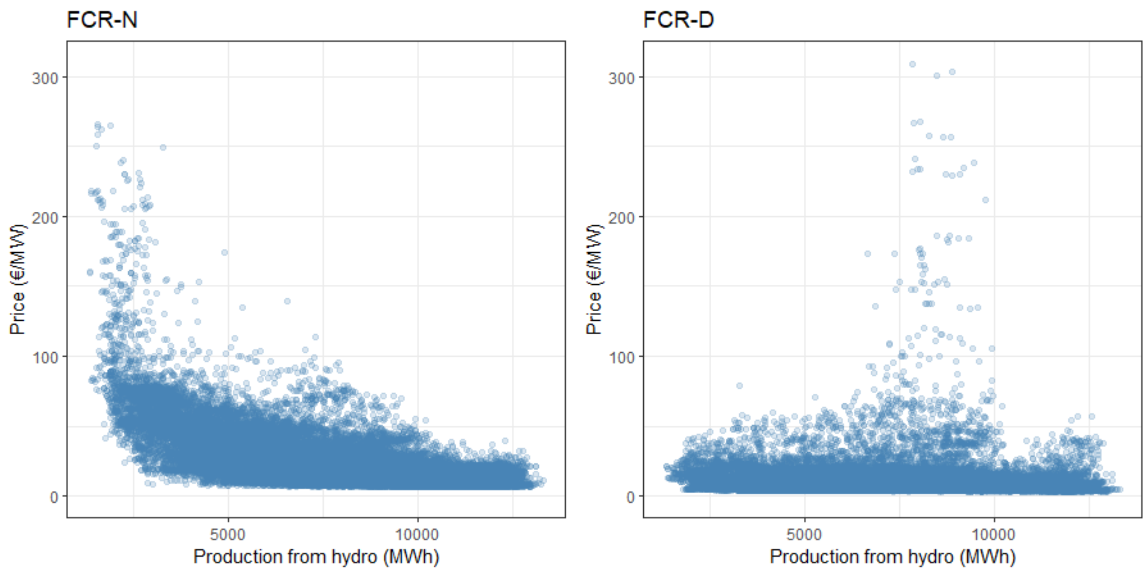


**Figure 3. FCR prices and hydro production**

It is not surprising that low hydro production is associated with high prices for FCR-N but not for FCR-D. A hydro resource that is operating at a very low set point can temporarily increase its output but not temporarily reduce it. It is therefore capable of providing FCR-D, but cannot provide FCR-N.

# Day-ahead prices

Another potential variable that might correlate with FCR prices is the day-ahead wholesale electricity price, which is related to the opportunity cost for a hydro producer to deviate from its otherwise optimal output level in order to provide FCR. However, this opportunity cost is also influenced by the internal water value for the hydro resource, which we do not observe. Both for FCR-N and FCR-D, a small positive correlation can be observed between the day-ahead price[1] and the FCR price, but the explanatory power of the day-ahead price variable alone is very low.

For the FCR-N price, it is interesting to note that a positive correlation with the day-ahead price is observed despite the negative correlation between the FCR-N price and hydro production shown above. Given that hydro production is positively correlated with the day-ahead price, the association between FCR-N prices and day-ahead prices is more positive when hydro production is controlled for. Nevertheless, the explanatory power of the day-ahead price remains relatively low.

# Reservoir levels

The amount of water in hydro reservoirs affect hydroelectric producers' ability and willingness to produce electricity, and thereby also their ability and willingness to offer FCR. We here look at how some indicators of hydro reservoir conditions correlate with FCR prices.

We use Swedish weekly aggregate hydro reservoir level data (measured in GWh electricity production equivalents). Since the aggregate reservoir levels follow a seasonal pattern, with high water levels in the late summer and early fall and low levels in the early spring, the reservoir level must be adjusted for seasonality in order to give a measure of whether the water levels are unusually high or low. We therefore construct a variable for the reservoir level as a fraction of the normal level for the given week.[2] We also construct a variable for the weekly change in reservoir levels, which gives an indication of whether there was a large inflow of water to the reservoirs at the time.[3]

Figure 4 shows how the FCR prices (as indicated by the color of the dots) have varied depending on our two indicators for the reservoir level conditions. Note that in this figure, since the reservoir level data is weekly, the FCR prices has been averaged by week. For FCR-D (the figure on the right), we can clearly see that the weeks with high average prices occurred when reservoir levels were quickly increasing and were far above normal for the season. This makes intuitive sense, since under such conditions the hydro producers likely wanted to produce at a high output level, making it costly to reduce output in order to provide FCR-D.

The story looks somewhat different for FCR-N (left side of Figure 4). As we have already seen, FCR-N tends to get more expensive if producers want to operate at a very low level (necessitating an increase in output in order to provide FCR-N). From Figure 4 it appears that this has occurred when reservoir levels were below normal but relatively stable or increasing.

---

[1] Here using day-ahead prices for SE2.

[2] The normal levels are here calculated as the average reservoir level for each week for January 2013 through June 2019.

[3] Change in reservoir level week t = ((reservoir level week t+1) – (reservoir level week t-1))/2.

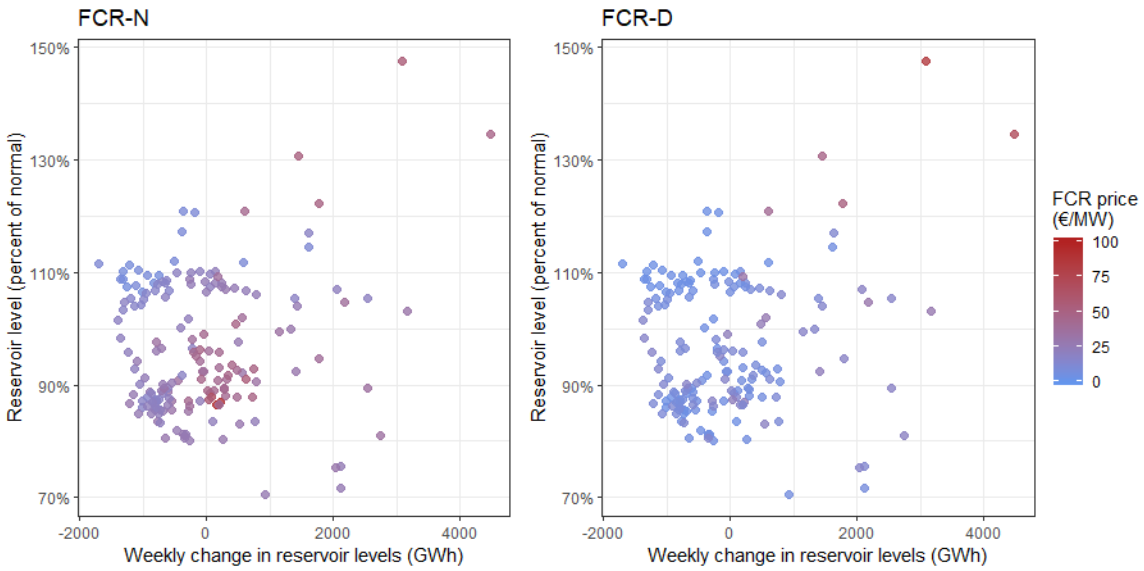


**Figure 4. FCR prices and reservoir levels**

## Demand for balancing services

The amount of FCR-N procured by Svenska Kraftnät has been essentially constant at 226 or 227 MW across all hours during the sample period. Variations in demand for FCR-N in Sweden can therefore not explain the price variations.

For FCR-D, the amount procured depends on the largest contingency, which in turn depends on which large power producers that are currently operating. Typically, the nuclear generator Oskarshamn 3 sets the requirement which means that about 427 MW of FCR-D is procured. When Oskarshamn 3 is not running, the procured amount may be somewhat lower. Nevertheless, the FCR-D procurement amount does not change much.

There is however another balancing service for which the demand could influence FCR prices. The aFRR (automatic Frequency Restauration Reserve) balancing service is a somewhat slower-responding balancing service compared to FCR and is activated with a signal from the TSO instead of with the frequency. But despite these differences, aFRR may be competing for the same resources that otherwise would supply FCR. High demand for aFRR could therefore spill over to higher FCR prices.

The amount of aFRR procured has varied considerably over the sample period. Svenska Kraftnät does not procure aFRR for all hours of the day. Further, the amount of aFRR procured has been gradually increasing over the sample period. For hours when aFRR was procured, Svenska Kraftnät typically procured between 100 and 200 MW in both the upward and downward directions. Figure 5 shows this variation for upward aFRR. The corresponding figure for downward aFRR looks very similar.

In a simple linear regression with the FCR price as outcome variable and the amount of upward aFRR volume as only explanatory variable, each MW of aFRR capacity is associated with an increase in the FCR price of about €0.04 (this is the case for both FCR-N and FCR-D). It is important to note that this should not be given a causal interpretation, given that the aFRR procurement is not random. However, we note that these estimates are only reduced slightly (to about €0.03) when fixed effects for hour, weekday, month and year are included in the regression.

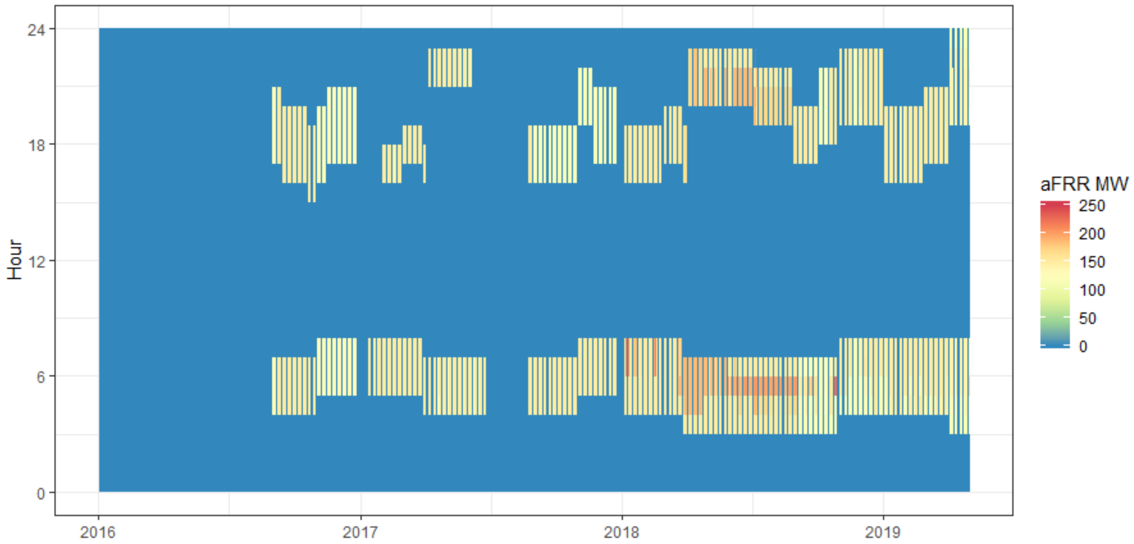


**Figure 5. Amount of upward aFRR procured**

# Linear regressions

We now combine the variables discussed above and analyze how much of the FCR price variation that can be captured by various linear regression model specifications. We do this by investigating how the $R^2$-values change depending on specification.

First, we consider how much of the variation in price that can be "explained" by only including fixed effects for hour, weekday and month. The first row of Table 1 shows the $R^2$ for these regressions. As seen, more than half of the variation in the FCR-N price is captured by these fixed effects, but only about 20% of the variation in the FCR-D price.

The second row of Table 1 shows the $R^2$ values for regressions using the weekly reservoir level data. This includes the reservoir level variable (both the absolute value and as a percent of normal for the week) and the variable for the weekly change in reservoir levels. Further, we construct a dummy variable for whether the weekly change is positive or negative and interact this with the reservoir level. The reason for including this interaction term is to allow for the effect of the reservoir level to be different between the inflow season and the depletion season. For FCR-D, these variables achieve an $R^2$ that is higher than that for the time fixed effects. However, the $R^2$ for the FCR-N regression is lower.

Next, we consider the group of variables that in Table 1 are referred to as "Market variables". This includes the hourly amount of electricity produced by Swedish hydro resources, the day-ahead market price and the amount of upward aFRR procured. It further includes the hourly Swedish electricity consumption, as well as production from wind and nuclear. Finally, because of the non-linear relationship between hydro production and FCR-N prices shown in Figure 3, we also include the square of hydro production as a variable. As seen in Table 3, these variables lead to an $R^2$ of about 60% for FCR-N, but less than 20% for FCR-D. For FCR-N, the hydro production variables alone give an $R^2$ of about 40% (not shown in the table).

When the reservoir and energy market variables are combined, the $R^2$ for FCR-N only increases slightly, indicating that the reservoir variables do not add much in terms of explanatory power once the energy market variables are included. For FCR-D however, the combined regression with both reservoir and energy market data gives an $R^2$ that is considerably higher.

Finally, adding the hour, weekday and month fixed effects increases the $R^2$ somewhat, especially for FCR-N. Nevertheless, it appears that the energy market and reservoir variables are capable of capturing most of the variation that can be captured by hour and time fixed effects. The energy market variables (especially hydro production) seem to be more important for FCR-N, and the reservoir variables more important for FCR-D.

**Table 1. $R^2$-values for various model specifications**

| | **FCR-N** | **FCR-D** |
|---|---|---|
| **Time fixed effects** | 0.523 | 0.216 |
| **Reservoir variables** | 0.199 | 0.287 |
| **Market variables** | 0.622 | 0.203 |
| **Reservoir & market variables** | 0.632 | 0.449 |
| **All of the above** | 0.710 | 0.484 |

# Predictive modelling

The above analysis fits linear regression models using all available data from the sample period. It is therefore possible that these regression models overfit the data and would not perform as well on new observations as the $R^2$-values in Table 1 would suggest. Further, due to non-linearities and interactions between the variables, it is possible that a more flexible model such as a neural network would deliver better predictions than what the above discussed linear regression models would.

Therefore, we here randomly split the data into a training set (60% of the hourly observations), a cross-validation set (20%) and a test set (20%). First, we explore how the linear regression models perform on the cross-validation set when the parameters are estimated using only the training set. Second, we perform a similar analysis using neural networks. Finally, using the test set data, we compare the predictive power of the best-performing linear regression model to the best-performing neural network model.

## Linear regression model predictions

Table 2 shows the root mean squared error (RMSE) on the cross-validation set for the same linear model specifications as discussed in the previous section. The model parameters were here estimated using only the training set, and the RMSE calculated when those parameters were used to make predictions on the cross-validation set. We here use RMSE as the evaluation metric instead of $R^2$, due to the interpretation difficulties for $R^2$ when applied out-of-sample.

The first row of Table 2 shows the RMSE for the baseline case where the average price from the test set is used to predict all observations in the cross-validation set. Table 2 shows that the models that include all variables give the best predictive performance, indicating that the models are not over-fitting the data too much. The prediction accuracy improves more for FCR-N than for FCR-D, compared to the baseline.

Energy Economics Lab

**Table 2. Cross-validation RMSE for various linear regression model specifications**

| | **FCR-N** | **FCR-D** |
|---|---|---|
| *Baseline* | 21.5 | 13.5 |
| **Time fixed effects** | 14.9 | 11.9 |
| **Reservoir variables** | 19.2 | 11.3 |
| **Market variables** | 13.0 | 12.1 |
| **Reservoir & market variables** | 12.8 | 10.1 |
| **All of the above** | 11.5 | 9.8 |

# Neural network model predictions

We now implement a neural network model for predicting FCR-N and FCR-D prices, using the same variables as above. We use a neural network architecture with one hidden layer containing 15 hidden units. Table 3 presents root mean square errors for the various groups of explanatory variables. Given that neural networks can capture non-linearities without the need for manually creating non-linear transformations of the input variables, we here do not explicitly include the squared hydro production variable or the reservoir level interacted with the inflow dummy.

As seen in Table 3, the neural network model results in lower prediction errors on the cross-validation set compared to the linear regression model. When using all of the explanatory variables, the RMSE is reduced by about 70% compared to the baseline for FCR-D and 50% for FCR-D.

In general, the RMSE are reduced when more variables are included, indicating that overfitting is not too much of a problem. It is interesting to note that, for both FCR-N and FCR-D, using all explanatory variables except the time fixed effects results in quite low prediction errors, meaning that the variables discussed above are sufficient to predict FCR prices with reasonable precision and capture a large amount of the price variation.

**Table 3. Cross-validation RMSE for various neural network model specifications**

| | **FCR-N** | **FCR-D** |
|---|---|---|
| *Baseline* | 21.5 | 13.5 |
| **Time fixed effects** | 14.4 | 12.3 |
| **Reservoir variables** | 16.5 | 8.0 |
| **Market variables** | 10.6 | 7.6 |
| **Reservoir & market variables** | 8.9 | 6.8 |
| **All of the above** | 6.4 | 6.8 |

# Comparison of modelling approaches

We conclude the analysis with a visualization of the predictions obtained from the linear regression models to those from the neural network model. For this we use data from the test set, and the estimated model parameters obtained when using the best-performing model for each case discussed above, i.e. the "All of the above" specification with all explanatory variables.

Figures 6 and 7 plot actual and predicted FCR prices for FCR-N and FCR-D, respectively. The predicted prices from the neural network model is shown on the left in each figure and the predicted prices from the linear regression model on the right.

For both FCR-N and FCR-D, we clearly see how the linear regression model fails to predict higher prices, especially those above 100 €/MW. The neural network models are much better at providing reasonable predictions for these cases. The neural networks are also better at predicting lower prices, but here the difference is perhaps less striking.

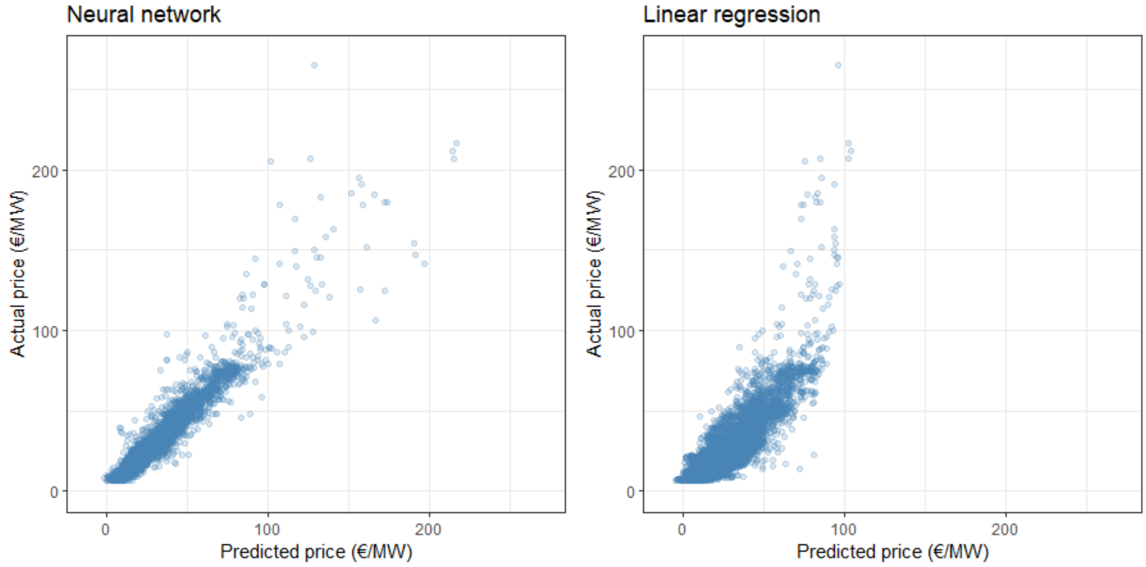


**Figure 6. FCR-N price prediction performance**

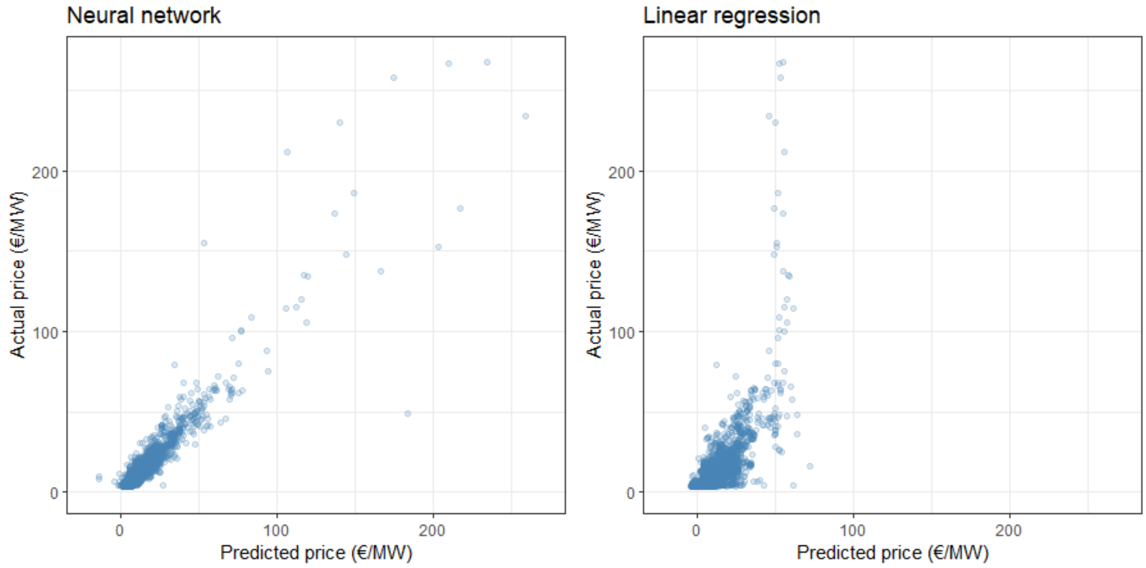


**Figure 7. FCR-D price prediction performance**

Energy Economics Lab is a center for research on energy markets and energy systems. We combine engineering and economics expertise in an interdisciplinary research environment to address market and business challenges related to the transition to a sustainable energy system. Energy Economics Lab is an initiative by RISE Research Institutes of Sweden.