My research

Nina Lange

March 24, 2025

My research centers around commodity and shipping markets. Initially, and still, my research dealt with term structure models for commodity prices and related markets such as weather or foreign exchange. Lately, several projects have focused on problems including both optimization and modeling of price input in order to support decision making in energy and shipping markets. The project presented in this short document deals with optimal operation of a large scale batteries as for instance Tesla Megapacks. The project is based on current work in progress with Stefan Røpke from DTU.

The decision problem

With the large increase in power production from renewable energy sources, volatility in electricity prices is high because of it's non-storable nature. A large scale battery connected to the power transmission grid can potentially help decrease the large variability in power prices and thereby support the transition to more green energy. The operator of an existing battery needs to decide 1) in which market(s) to trade and 2) when to trade and how much to trade. We assume that the operator of the large scale battery can trade directly in the wholesale electricity markets and that all bid and offers are matched[1]. We also assume that the size of the battery and the trades are not large enough to affect the market prices or behavior of other participants.

Regarding 1), the operator has – in theory[2] – multiple electrivity markets to trade in. They could trade electricity futures or forwards. This is in most cases less interesting as most of these derivatives would be a product requiring the parties to exchange electricity for a pre-fixed price at a steady rate over the day, whereas the point of these large scale batteries is profit from hourly price differences. The most relevant markets to consider are the day-ahead markets and the intraday market, where electricity are traded at an energy exchange for delivery in a specific time interval in the same or next day. For the day-ahead markets, price and volumes are submitted by a certain deadline (referred to as gate closure) after which clearing prices (often hourly or half-hourly) are determined and order volumes are reported back to market participants. In intraday markets, electricity are traded for specific (short) time intervals in a continuous markets up until shortly before the actual time interval (which usually spans 5-60 minutes). Market participants can submit bids and offers, and prices are formed based on completed transactions. Finally, as supply and demand must balance at all time, the balancing market takes over. Here, a system operator ensures the balance of supply and demand by reducing or increasing the amount of electricity feeded into the grid.

For the empirical work in this project, we consider the UK electricity market. Here, an operator can trade in five different auction type markets with four different gate closures. We refrain from including the continuous-type intraday markets due to data availability as well as problem complexity.

[1] In reality, a market participant should report both the amount of MWh, she wants to trade and at which price.

[2] Not all market participants have access to e.g., the balancing market or intraday markets and the electricity derivatives market is of less interest to a battery operator.

Regarding 2), the operator has to schedule the operation of the battery according to some objective function. This could be stated as a stochastic programming problem, an objective that balances profit and potential losses, or, as in our case, a model that computes a plan that maximizes (expected) profit over a two-day horizon. The plan is then fixed for one market at a time and subsequently adjusted when new market information arise. The objective function maximizes profit over all markets by choosing the amount $x_{m,h}$ to sell or buy in each market for each time-period.

$$\max_{x} \sum_{h \in H} \sum_{m \in M} p_{m,h} x_{m,h} \tag{1}$$

The trading decision is made subject to the number of constraints that ensures that the plan is feasible. The battery has a certain capacity and charging rates, and the number of cycles is limited in order to not wear out the battery. Also, constraint are added such that the battery is not emptied at the end of the first day. The base model allows speculative trading where a buy order in one market is canceled out by a sell order in another market. To disallow this, we can enforce that all trades in a time period should be in the same direction (buy/sell) by adding constraints.

Evidently, the electricity prices are not known at the time of the decision, but must be forecasted prior to gate closure using an appropriate model and available market information. The model runs shortly before each gate closure and includes previously agreed trades and updated price forecasts. The first two markets close at the same time in the morning before the day of the delivery of electricity. These two markets trade hourly contracts for the following day. The third and fourth market closes in early resp. late afternoon for half-hourly contracts in the following day. The fifth market closes the next morning and it is only possible to trade in half-hourly contracts for the last half of the day.

Input to optimization model

As model input, we use a Tesla megapack for model parameters. We explore different reduced form models such as the Ornstein-Uhlenbeck process together with a number of machine learning models and simple moving average forecasts.

When forecasting electricity prices for the optimization model, we forecast a panel of data. The panel of actual prices are revealed shortly after making a bid to the market. In the literature on electricity price forecasting, the performance of a model is evaluated using classical metrics such as RMSE or MAPE. Although these are relevant for the purpose of evaluating the performance of forecasts before entering the forecasts into the optimization model, it might not be the best metric to ensure optimal use of the battery. For example, if RMSE is zero, the model is perfect, and the battery is operated optimally. But the opposite is not true: If the battery runs optimally, RMSE is not necessarily close to zero and these metrics are further unable to decide between forecasting models. Consider, for instance, a simple example with a battery that is only charged and discharged one half hour per day. In this case optimal operation only requires predicting when the the highest and lowest price occur[3]. If the model predicts 200 as the highest half-hour price and 10 as the lowest half-hour price and the realized values are 55 resp. 45, RMSE would be high, but as long as the most and the least expensive half-hour are correctly chosen, the forecast should be considered perfect in terms of extracting profit from battery operation.

[3]And if the spread is large enough to account for efficiency losses.

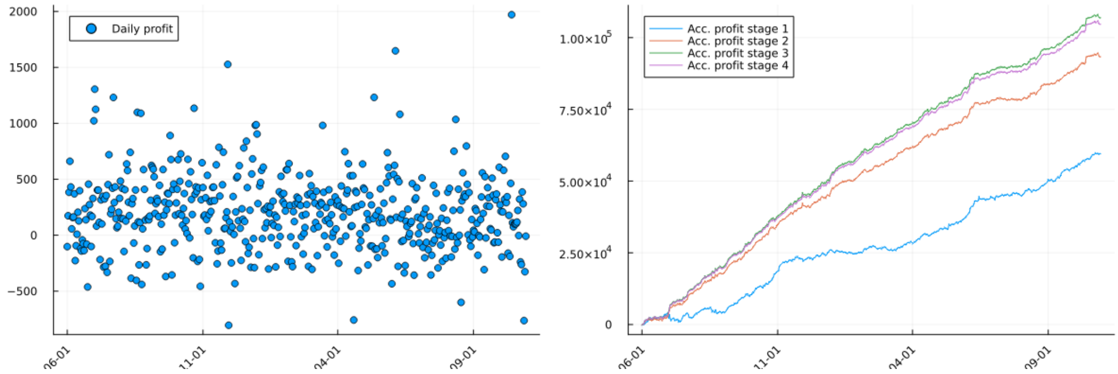

Figure 1: Daily profits and accumulated profits

For optimal charging and discharging, the more important task is to predict the highest and lowest values and when they occur, which makes it natural to consider measures that measure the ability of a forecast to rank the hourly/half-hourly prices correct. We discuss and calculate measures like Kendall's tau, which measures the correlation of ordering of values and precision@k/precision@k-of-m that measures the ability to forecast the highest/lowest values and especially when these are not sufficient to measure the accuracy of price forecasts for the purpose of battery operation.

We therefore consider a metric that in essence measures the ratio between the profit from buying in the anticipated k cheapest half hours and selling in the anticipated k most expensive half hours and the profit from operating in the realized best half hours. For example, if a battery charges (and discharges) in two hours, the price forecast should correctly identify the two hours or four half-hours with the highest prices and the two hours or four half-hours with the lowest prices. It's not a perfect measure for our use, as the physical constraints could limit us from actually charging and discharging at perfect hours in terms of profit. The measure does not consider the ability to capture the order, but only the sum of the hours the model forecast to be best relative to the actual best hours.

Output of model

For the best performing forecasts, we analyze how the model perform with and without speculation and based on the frequency of rerunning the model and thereby including new information available before gate closure. An example is presented in the figure above. As can be seen on the left figure, daily profits are highly volatile, but accumulated profits are quite stable over time. In the specific case, the profit increases quite substantially when the schedule for later markets are revised as new information prevails. The addition of the fifth market (that closes in the morning in the day of trading) does not add much value. This is not a huge surprise, as it only covers that last half of the day and when the plan is adjusted prior to gate closures, earlier made obligations to buy and sell leaves little room for big changes given the battery constraints.