FACULTY OF SCIENCE
UNIVERSITY OF COPENHAGEN

Bachelor Thesis

Björn Pettersson zts900@alumni.ku.dk [Code]

Aggregated FCR-D participation for BTM energy communities

Supervisor: Thomas Troels Hildebrandt

Jun 2026

Aggregated FCR-D participation for BTM energy communities

Björn Pettersson [1]

Abstract

With the transition towards more renewable energy resources and a more distributed energy grid, the demand for ancillary services is projected to rise. While previous research has looked at potential participation of citizen energy communities (CEC) (Crowley et al., 2024), the business case in use today in Denmark is that of the Behind-the-Meter (BTM) energy community. This study looks at the feasibility for BTM communities with solar PV and batteries to participate in the ancillary services markets through a Balancing Service Provider (BSP) aggregator for the DK2 market in the Nordic frequency grid. A case study is conducted, studying this type of aggregator in order to gain access to real data and insight into the practical case. For the BSP, the available ancillary services products Frequency Containment Reserve – Disturbance (FCR-D) up and down seem most promising. Since FCR-D bidding is done through two day-ahead auctions, early and late, we derive two mixed integer linear programs (MILP) to find the optimal charge-discharge schedule for an aggregated BTM battery as well as the optimal bid volumes for FCR-D up and down, given forecasted data. Here objectives of self consumption and regular battery arbitrage compete with the goal of reserving capacity for ancillary services payments, which locks up battery capacity and hence restricts state of charge (SOC) movements for other purposes. Given an upper bound based on feeding perfect forecasts into the program for 2025, we observe that stacking FCR-D did not significantly reduce self consumption revenue and provides significant upside. This result seem to persist through a number of different portfolio configurations. Furthermore, conducting a forecast noise sensitivity analysis we find that a simple seven day persistence forecast can yield 88 percent of the upper bound returns. This suggests that the optimal strategy when stacking FCR-D up and down may be quite simple and relatively robust to forecast noise as long as the order of magnitude between values stays roughly the same.

1. Introduction

1.1. Energy communities today

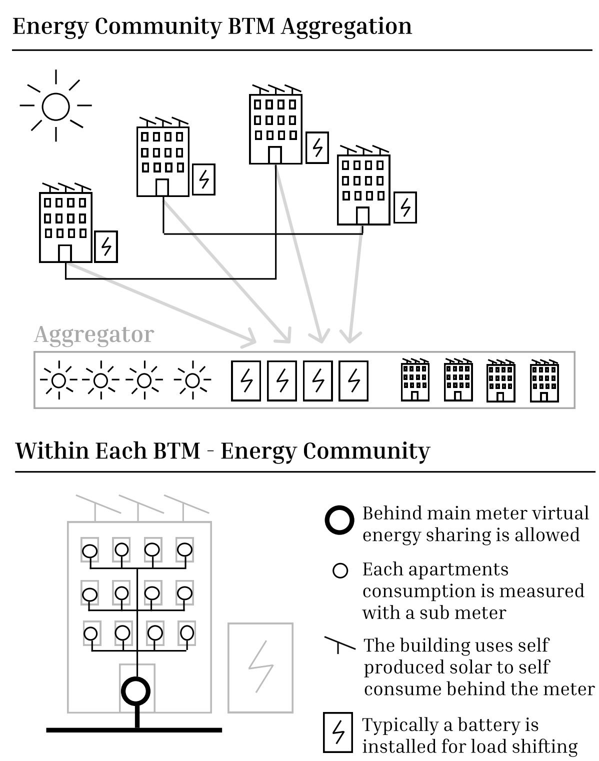

Figure 1. BTM energy community and aggregation, an aggregator can virtually view these energy assets as one and use that virtual asset to bid into ancillary services markets

With ambitious climate targets of neutrality by 2050 (Committee on Climate Change, 2020), the EU needs quick climate adaptation. One of the main culprits is building-related systems that account for 30 percent of final energy consumption and 26 percent of global energy-related carbon emissions (International Energy Agency, 2023). One promising strategy is to move towards a more distributed energy grid where on- or near-site RE production and batteries can offload the grid and reduce the net energy demand of the building; this has been shown to be an effective approach

(Amini Toosi et al., 2022). With the REDII (European Union, 2018) and IMED (European Union, 2019) directives, the concept of energy communities is brought forward as a concrete way of transitioning towards a more distributed energy grid, but while member states are obligated to implement the concepts, the specifics and market incentives around them remain vastly different across member states. In Denmark, while talks about behind-the-transformer style energy sharing or sharing in larger areas such as Citizen Energy Communities (CEC) have been ongoing, the practical business case in use today is for on-site energy sharing, Behind-the-Meter (BTM), one such energy community is Vendsysselhus (Enyday ApS, 2024). We also see similar concepts across Sweden with *gemensam el*, and Germany's *Mieterstrom*, and this seems to be a first step towards larger energy sharing in northern Europe (REScoop.eu, 2024). Today the business case of BTM energy sharing is not tied to the legal EU definitions of CEC's or renewable energy communities REC's, but we can expect closer adoption in the future. Hence the relevance of energy communities and distributed resources will only grow.

The main business case for a BTM, see figure 1 setup today consists of substituting the building's energy consumption with on-site produced solar PV, which is cheaper and, furthermore, since it does not enter the district energy grid, it is tax and tariff free. The reason CEC setups remain unfeasible today is because energy sharing between buildings would require energy entering the district grid, where tax and tariffs, which is a very large part of the final consumer price, would be applied. Since the sell price is much lower than the purchase price because of tariffs applied to purchase but not to sell price, producing energy in one building and consuming it in the building next to it would be done at a loss, at least for a residential PV producer. This leaves the BTM case as the most feasible today. These setups often consist of a larger multi-unit housing block, where solar PV and batteries are installed. This can also be combined with heat pumps, smart appliances or charging of electric vehicles (EVs). As the price of solar and batteries both are shrinking, the demand for these types of setups also seems set to rise. It is this business case that is driving the transition at the moment towards a flexible grid, hence further improvements of the case through additional income streams with ancillary services would further speed things up.

1.2. Ancillary services

Because many renewable energy (RE) sources are weather and climate dependent, take solar and wind. Energy production from these sources is more uncertain, and especially in the case of solar PV, production will be concentrated during certain daylight hours, which is not necessarily when we want to consume the most; this is typically towards the evening, when people are making dinner and arriving home from work. As we build out these resources to replace more stable gas or coal turbines with predictable and adjustable output, our grid stability becomes more uncertain. The operator of the high voltage energy grid (TSO), in Denmark Energinet, has to make sure the grid frequency stays stable and does not cause blackouts if it drops too low, meaning too high demand, or equipment damage if it rises too high, meaning too much production. It does this by buying buffer capacity through ancillary services. It pays operators of energy resources, amongst others batteries, to be ready to either put more energy out into the grid (up regulation), or to increase consumption (down regulation), in case the frequency moves beyond a certain band around the target frequency of 50 Hz.

1.3. FCR-D bidding

FCR-D stands for Frequency Containment Reserve Disturbance. It is a specialised, fast-acting grid stabilisation service that transmission system operators (TSOs like Energinet) buy to prevent blackouts when something suddenly goes wrong on the power network. There is one version for up regulation, putting energy into the grid FCR-D up, and one for down regulation taking energy out of the grid, FCR-D down. The disturbance part refers to FCR-D being activated when the grid frequency suddenly falls out of the normal range, around 50 Hz. Batteries are well suited for FCR-D, since it requires a fast response time.

The bidding process is outlined in the implementation guide (Energinet et al., 2026), in Figure 2. In this study we will use forecasts and a linear solver to first forecast what we need to know, and given what we know we will use the linear solver to solve for optimal bidding and scheduling. These forecasts and bids will be submitted within the outlined process. The time line for the two auctions as well as early and late forecasts is shown in Figure 2.

The bidding process consists of two auctions. The TSO is looking to fill its demand for FCR-D down and FCR-D up each hour; bidders that have energy resources for either up or down regulation submit bids for reserving capacity that will be activated automatically in the case of a frequency deviation. For each hour bids are submitted through the Nordic MMS bid portal. This program is a simplification since it assumes all bids are accepted, but in reality there can be partial acceptance of bids. The program also derives the optimal hourly bids, while in practice Nordic MMS takes block bids but also allows several bid submissions to stack hourly bids, which may be somewhat more complicated. For battery assets, a Normal Energy Management (NEM) mode can be activated that provides some support in regards to returning to a desired SOC after FCR-D activation; this is also omitted from the model. We do however keep some capacity reserved to make sure we can deliver on our reservation commitments (Energinet, 2023).

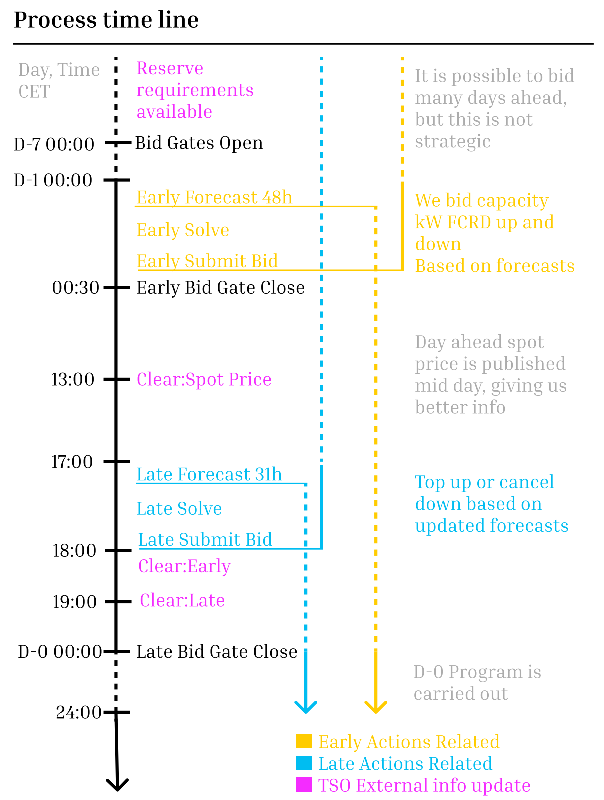

Figure 2. Time line for early and late auction, as well as bids. We want to forecast as late as possible, solve the program giving us optimal reservations for FCR-D up and down as well as optimal charge schedule. We commit bids before bidding gates close. The late auction uses accepted early bids and the cleared spot price from 13:00, in addition to the forecasted variables. Forecasts are made 48 and 31 hours ahead but only the last 24h of the following day is used.

1.4. Literature Review

Rehman et al. (ur Rehman et al., 2025) provide a recent systematic review of energy management models for smart homes covering 2018–2024, confirming that mathematical programming approaches – particularly mixed-integer linear programs (MILP) fed with forecasted consumption, production, and price data – remain the dominant methodology for optimal dispatch and scheduling of distributed energy resources. Their bibliometric analysis identifies ancillary service participation and AI/ML-driven scheduling algorithms as emerging but underexplored research areas, motivating the present study's focus on stacking FCR-D bidding with self-consumption optimisation for BTM energy communities. Williams (Williams, 2013) provides foundational guidance on the formulation of such programs, including the treatment of minimum bid constraints through binary variables, which introduces the mixed-integer structure central to our model.

Closest to the present work, Crowley et al. (Crowley et al., 2024) propose a two-stage stochastic program for an energy community bidding into the Nordic ancillary service markets. Their model considers multiple asset types (inflexible consumers, shiftable loads, and battery prosumers) and simultaneous bidding across FFR, FCR-D, FCR-N, and mFRR markets in the DK2 area. Their key finding is that the community allocates bids almost exclusively to the FCR-D and FCR-N markets, with negligible FFR participation and only marginal mFRR bidding from shiftable prosumers. An out-of-sample analysis shows that 100 in-sample scenarios suffice for reliable service delivery in over 96% of cases. The present study adapts and extends Crowley et al.'s formulation for the behind-the-meter (BTM) energy community case, applying the power balance constraint at the building level rather than the aggregate community level, focusing exclusively on FCR-D (the most accessible product for a BSP in DK2), and introducing sequential early and late auction programs to reflect the actual market timeline.

The sensitivity of optimal scheduling to forecast accuracy is examined by Langtry et al. (Langtry et al., 2024), who study model predictive control (MPC) for building energy systems with distributed generation and storage. They introduce synthetic forecast noise using a Gaussian random walk model and show that MPC performance is most sensitive to electricity price and carbon intensity forecast errors, while being relatively robust to solar generation prediction inaccuracy. Their methodology of scaling noise relative to empirical persistence forecast error provides a good framework for sensitivity analysis. This study adopts a similar noise model to investigate how revenue from the MILP-based scheduling degrades under forecast uncertainty.

Another important reference is Thingvad et al. (Thingvad et al., 2024), who define detailed constraints for Limited Energy Reservoir (LER) type assets, covering many battery storage systems (BESS). They derive a program incorporating both arbitrage and ancillary services participation in FCR-D up, FCR-D down, FCR-N, and FFR. The model structure is similar in terms of the MILP base, and we later use these LER constraints to extend our own model. Formally, LER assets are those that cannot sustain full energy activation for four consecutive hours. For FCR-D, LER prequalification requires reserving energy equivalent to 20 minutes of full output ($\upsilon_E^D = 0.33\,\text{h}$ per MW) (Energinet, 2023), as well as 20% of the FCR-D capacity in the opposite direction for NEM recharging after activation. What Thingvad et al. do not consider is the inclusion of the second auction or the BTM type asset aggregation 1 with building-level energy constraints. The derivation of how the LER constraints are integrated into our model is given in Appendix A, Steps 8–9.

1.5. Scope selection

For the purpose of this study, the data obtained for the case study is based on DK2 and the Nordic grid. The legal status of the aggregator will be assumed as Balancing Service Provider (BSP), as this is vastly more accessible for market entry than the Balance Responsible Party (BRP) status that requires collateral of 1 million DKK compared to 100k DKK for the BSP registration. The BSP status only allows participation in markets that provide reservation payments only, meaning you get paid for reserving capacity but not for the energy provided if that capacity is activated. This is typical for products with faster activations such as FFR, FCR, FCR-D and FCR-N, compared to products that are slower but require a higher sustained energy output, such as aFRR and mFRR.

Based on these restrictions only a few instruments are available for a BSP in DK2: FFR, FCR-D up and FCR-D down.

The paper "Can Energy Communities Provide Ancillary Services? A Stochastic Program for Aggregated Bidding in the Nordic Ancillary Markets" (Crowley et al., 2024) provides a linear program for optimal bid quantities for a CEC energy community. The program is stochastic where each day during the study period helps build a profile of possible outcomes and probabilities of these outcomes. Based on this stochastic daily profile Crowley et al. derive optimal bid quantities for the DK2 instruments FFR, FCR-D up, FCR-N, mFRR and aFRR, deriving a 24h optimal bid vector. They find that the optimal program chooses to bid into FCR-D and FCR-N, with approximately a 50/50 split, and placing almost no bids for the very fast FFR.

It is important to point out that the BTM asset type we are looking at comes with batteries and solar preinstalled on the basis that self consumption and arbitrage will give a reasonable payback period. Since reserving capacity for ancillary services also requires that we are ready to deliver this capacity if necessary, this means that the goal of reserving as much capacity as possible is restricting how much of the battery can be used for self consumption and arbitrage at any given hour. In this sense these objectives are slightly competing with each other. But furthermore, the implication is that any revenue increase from the base case with only self consumption and arbitrage is attractive, since the battery is not mainly financed on the basis of ancillary services revenue.

While FCR-D up, FCR-D down and FCR-N bidding is done through the Nordic MMS portal, bidding for FFR is separate. For the sake of this study we have hence focused on the FCR-D up and FCR-D down products, as per the Crowley et al. study these are believed to be vastly more attractive than FFR.

Primarily we are interested in the financial incentives for the BTM type asset to participate in the ancillary services market: how attractive is it, and how does adding this additional objective of FCR-D up and down participation on top of existing self consumption and arbitrage objectives work in practice? Does the restriction of battery SOC significantly reduce self consumption revenue or can these two goals exist harmoniously together?

2. Data handling

The general use of data within this study was to contextualise and back test the derived model in order to evaluate performance. All data is for the year 2025. See Figure 3 for overview and Figure 5 for how data is fed to the model. For an overview of the data sources see Table 1. For actual references to column names in the dataset as well as units see the appendix Table 10, table 4 shows graphically how variables vary across days.

2.1. Energy community data, Production and Demand

Energy community data obtained through industry collaboration with Danish energy community manager Enyday (Enyday A/S, 2025) was used to construct an artificial aggregate portfolio which consists of a set of $E = 10$ energy communities, in order to meet the minimum bid size for FCR-D of 0.1 MW; the portfolio has capacity 1 MW. The obtained datasets consist of all energy meters from apartments within larger residential buildings with solar PV; the data is cumulative and contained some null values which were linearly interpolated. The cumulative nature ensures that the total consumption remains realistic. We then convert this to a community aggregate per-hour consumption and production. We generally see a consumption spike when people arrive home from work 17:00–20:00. PV production generally has a gradual peak during the day and is very seasonal, with almost none in winter. In the model we observe the data in $PV_{e,t}$ and $D_{e,t}$.

2.2. Energy Price Data, Import and Export Price

We construct the prices for importing and exporting electricity for the energy communities to the local distribution grid, owned by the distribution system operator DSO. The spot price is a variable price set the day ahead at NordPool P^{spot}. These are obtained in two datasets, since NordPool switched from hourly frequency (Energinet, 2025c) of measurement to per 15-minute basis for the second half of the year. We average these prices per hour (Energinet, 2025b). On top of that the DSO Radius adds tariffs T^{DSO} that vary depending on what hour in the day it is and what season in the year (Radius Elnet A/S, 2025). F^{state} are the state-related expenses for transport and system fees. On top of that a VAT is added μ_{VAT}. A feed-in tariff is paid for using the grid to sell electricity $C_{\text{feed-in}}$.

$$P_t^{\text{im}} = \left(P_t^{\text{spot}} + T_t^{\text{DSO}} + F^{\text{state}}\right) \times \mu_{\text{VAT}} \tag{1}$$

$$P_t^{\text{ex}} = P_t^{\text{spot}} - C_{\text{feed-in}} \tag{2}$$

2.3. FCR-D Data, prices and activations

Ancillary services related data used is the grid frequency to calculate when and how much FCR-D participation affects state of charge SOC. Since the Nordic frequency grid consists of both DK2, Sweden and Finland, we use data from the Finnish TSO on grid frequency (Fingrid, 2025). The data resolution is per 0.1 seconds, and we convert these to hourly fractions where the frequency left the band. The activation shares are pre-computed from 0.1-second resolution Nordic frequency data in the model $y_t^{\text{FCR-D}\uparrow\text{act}}$ and $y_t^{\text{FCR-D}\downarrow\text{act}}$. These fractions are typically small for FCR-D and are included mostly for correctness but have little impact on model results in practice.

In each auction the TSO wants to fill some demand for ancillary services each hour; it pays for these by a price-as-cleared mechanism, meaning everyone gets the price of the highest accepted bid. We obtain these historical prices for 2025 from the Energi Data Service (Energinet, 2025a). Given that your bid is accepted fully this is how much you get paid for reserving capacity in a given hour, in the model $\lambda_t^{\text{FCR-D}\downarrow\text{early}}$, $\lambda_t^{\text{FCR-D}\downarrow\text{late}}$, $\lambda_t^{\text{FCR-D}\uparrow\text{early}}$, $\lambda_t^{\text{FCR-D}\uparrow\text{late}}$.

3. FCR-D Early Formal Definition

In this section we present the complete formal definition of the mixed-integer linear program (MILP) used to determine the optimal aggregated bidding schedule in the FCR-D early auction (closing 00:30 the day before delivery), jointly with self consumption and arbitrage. The model is derived from the FCR-D base case of Crowley et al. (2024) with several modifications specific to the BTM aggregation setting; the full derivation is given in Appendix A and a summary of where equations stem from can be found in appendix table D. The complete listing of equations can be found in appendix section C.

3.1. Variable definitions

Table 2 contains the variable definitions for the early model. At the top we see the sets we loop over when solving: the set of time steps T, and since we have some energy community specific parameters and constraints such as the building

Table 1. Model input parameters and their data sources.

Symbol	Description	Source
Energy community Data		
$D_{e,t}$	Hourly demand	(Enyday A/S, 2025)
$PV_{e,t}$	Hourly PV production	(Enyday A/S, 2025)
Energy Price Data		
P_t^{spot}	Spot price DK2	(Energinet, 2025c;b)
T_t^{DSO}	DSO tariff	(Radius Elnet A/S, 2025)
λ_t^{im}	Import price Eq. (1)	(Energinet, 2025c;b; Radius Elnet A/S, 2025)
λ_t^{ex}	Export price Eq. (2)	(Energinet, 2025c;b)
FCR-D Data		
$\lambda_{\text{early},t}^{\text{FCR-D}\downarrow}$	FCR-D down early price	(Energinet, 2025a)
$\lambda_{\text{late},t}^{\text{FCR-D}\downarrow}$	FCR-D down late price	(Energinet, 2025a)
$\lambda_{\text{early},t}^{\text{FCR-D}\uparrow}$	FCR-D up early price	(Energinet, 2025a)
$\lambda_{\text{late},t}^{\text{FCR-D}\uparrow}$	FCR-D up late price	(Energinet, 2025a)
$y_t^{\text{act}\uparrow}$	FCR-D up activation	(Fingrid, 2025)
$y_t^{\text{act}\downarrow}$	FCR-D down activation	(Fingrid, 2025)

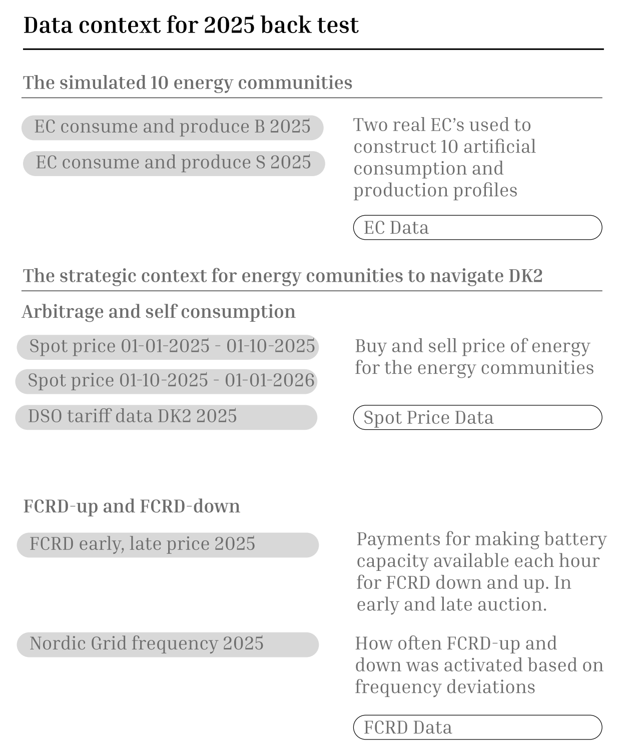

Figure 3. Overview of types of input data, and their purpose within the model. The purpose-based combined data sets are represented with an outline for diagrammatic use in Figure 5.

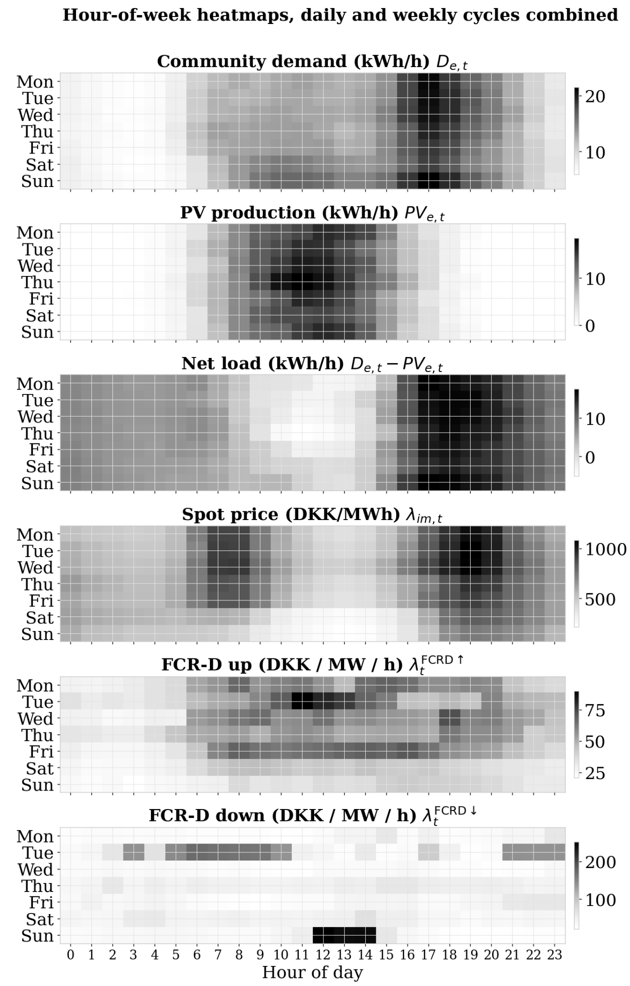

Figure 4. Overview of averaged weekly values for the datasets. We observe the cyclical daily nature of demand and PV production as well as spot price, with a weekend effect. We also observe that FCR-D prices show much less daily cyclical predictability comparatively.

power constraint, we also loop over the set of energy communities E.

The table distinguishes between model parameters, which are either scalars that define our case – here for each energy community e's battery its max power, state of charge and efficiency $\{\bar{b}_e, \bar{S}_e, \eta_e\}$ – or parameters governing the rules regarding how much of our bidded capacity we need to physically reserve in our battery state of charge through the Normal Energy Management (NEM) requirements: the energy headroom and power headroom required $\{\upsilon_E^D, \upsilon_P^D\}$, as well as the minimum and maximum bid size $\{p^{min}, \bar{P}\}$.

It also includes vectors that vary across e and t. They are forecasted based on 2025 data; for this first study we construct an upper bound using perfect forecasts. These variables include price variables denoted with λ: for arbitrage-related import and export energy prices $\{\lambda_t^{im}, \lambda_t^{ex}\}$, and for the reservation prices for FCR-D down and FCR-D up $\{\lambda_t^{FCRD\uparrow}, \lambda_t^{FCRD\downarrow}\}$. In addition there are energy community specific variables, namely how much energy is consumed and produced with solar PV $\{D_{e,t}, PV_{e,t}\}$.

Furthermore there are vectors over t that are related to how the bids are processed by the buyer of these ancillary services, the transmission system operator (TSO): namely whether reserved capacity is activated and in practice used to stabilise the grid $\{y_t^{actFCRD\uparrow}, y_t^{actFCRD\downarrow}\}$, and whether bids are accepted $\{y_t^{accFCRD\uparrow} = 1, y_t^{accFCRD\downarrow} = 1\}$. Because of lack of data and the fact that we bid at price zero, we assume all bids are accepted.

The table also contains variables solved for internally by the model when optimising the constraints. Mainly what we obtain after solving is the optimal charge/discharge schedule through $\{b_{e,t}^{ch}, b_{e,t}^{dis}\}$ as well as the optimal bid vectors of how much capacity we want to reserve each hour for each energy community, $\{p_{e,t}^{FCRD\uparrow res}, p_{e,t}^{FCRD\downarrow res}\}$. Obtaining these variables also determines how much we import and

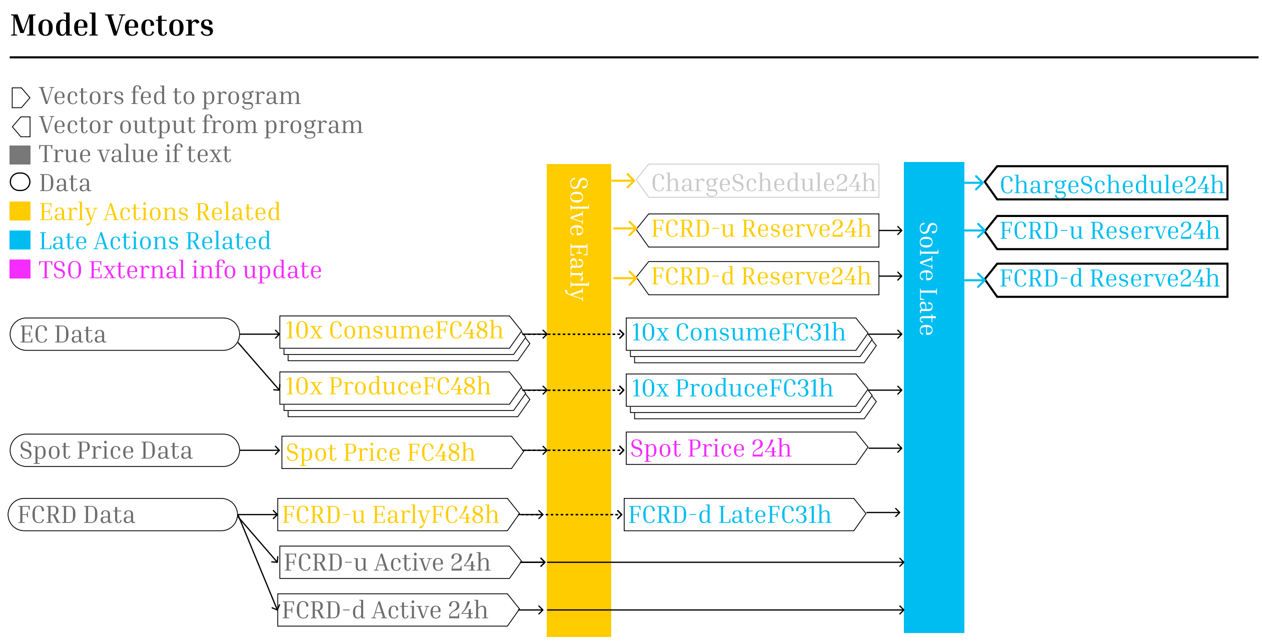

Figure 5. The first auction is held to obtain an initial schedule and bid vectors, the second to top up or cancel down the early bids. The figure shows how vectors are fed into the different models and from what data they stem. Forecasted vectors are denoted FC with the length of the forecast. Where a forecast vector is longer than 24h, the last 24h is used as this corresponds to the day the program is run. The programs are similar but the late program takes in the accepted reserve prices output from the early auction. For the sake of this model we assume all bids are fully accepted as we will set bid price 0.

export as well as net consumption $\{p^{im}_{e,t}, p^{ex}_{e,t}, p_{e,t}\}$, and the battery's state of charge $SOC_{e,t}$.

We include how much of our reserved energy is actually activated and used to stabilise the grid $\{p^{FCRD\uparrow act}_{e,t}, p^{FCRD\downarrow act}_{e,t}\}$, and finally binary variables that are used to encode the non-linearity stemming from our minimum bid constraint – we can either bid 0, or within the range $[p^{min}, \bar{P}]$ – they are $\{z^{\uparrow}_t, z^{\downarrow}_t\}$.

3.2. Objective Function

Derived from Crowley et al. (2024) (Eq. C-2a), with FCR-D down revenue added. The objective minimises the net cost of operating the aggregated portfolio over the day. We minimise what we spend on imported energy C^{im} minus export revenue R^{ex}; this implicitly maximises self-consumption and arbitrage. Adding FCR-D up and down, we simply subtract (since we are minimising) the revenue obtained from participating in these markets $R^{FCRD\uparrow}$, $R^{FCRD\downarrow}$:

$$\min \sum_{t\in T} \left[C^{im}_t - R^{ex}_t - R^{FCRD\uparrow}_t - R^{FCRD\downarrow}_t \right] \quad \text{(E-obj)}$$

The cost of importing C^{im} is the import energy price λ^{im}_t multiplied with the aggregated energy imported $\sum_{e\in E} p^{im}_{e,t} - \sum_{e\in E} p^{FCRD\downarrow act}_{e,t}$. We do not pay or get paid for the FCR-D related imports $\sum_{e\in E} p^{FCRD\downarrow act}_{e,t}$ or exports $\sum_{e\in E} p^{FCRD\uparrow act}_{e,t}$; FCR-D payments are made purely through the separate revenue terms $R^{FCRD\uparrow}$, $R^{FCRD\downarrow}$, hence we subtract these from the import and export volumes respectively. The export revenue has a similar structure:

$$C^{im}_t = \lambda^{im}_t \left(\sum_{e\in E} p^{im}_{e,t} - \sum_{e\in E} p^{FCRD\downarrow act}_{e,t} \right) \quad (3)$$

$$R^{ex}_t = \lambda^{ex}_t \left(\sum_{e\in E} p^{ex}_{e,t} - \sum_{e\in E} p^{FCRD\uparrow act}_{e,t} \right) \quad (4)$$

The FCR-D up and down revenue terms $R^{FCRD\uparrow}$, $R^{FCRD\downarrow}$ are similarly the reservation price for each instrument $\lambda^{FCRD\uparrow}_t$, $\lambda^{FCRD\downarrow}_t$ multiplied with the aggregated FCR-D reservation for each instrument: up $\sum_{e\in E} p^{FCRD\uparrow res}_{e,t}$ and down $\sum_{e\in E} p^{FCRD\downarrow res}_{e,t}$:

Table 2. Sets, parameters and decision variables for the early FCR-D model.

Symbol	Description
Sets	
E	Set of battery-prosumer energy communities
T	Set of hourly time steps in the day
Parameters	
λ_t^{im}	Import price for hour t (DKK/kWh)
λ_t^{ex}	Export price for hour t (DKK/kWh)
$\lambda_t^{FCRD\uparrow}$	FCR-D up reservation price, hour t
$\lambda_t^{FCRD\downarrow}$	FCR-D down reservation price, hour t
$PV_{e,t}$	Forecasted PV production for e at t
$D_{e,t}$	Forecasted demand for e at t
$\bar{b}_e$	Max battery power for e (kW)
$\bar{S}_e$	Max battery energy capacity for e (kWh)
η_e	One-way battery efficiency for e
$y_t^{actFCRD\uparrow}$	Pre-computed up activation share, hour t
$y_t^{actFCRD\downarrow}$	Pre-computed down activation share, hour t
$y_t^{accFCRD\uparrow}$	Up bid acceptance indicator (= 1)
$y_t^{accFCRD\downarrow}$	Down bid acceptance indicator (= 1)
υ_E^D	LER sustain duration (= 1/3 h)
υ_P^D	LER NEM power reservation factor (= 0.20)
p^{min}	Minimum portfolio bid size (kW)
$\bar{P}$	Maximum portfolio bid size (kW)
Decision variables	
$p_{e,t}$	Net consumption of e at t
$p_{e,t}^{im}$	Grid import for e at t (≥ 0)
$p_{e,t}^{ex}$	Grid export for e at t (≥ 0)
$b_{e,t}^{ch}$	Battery charging power, e at t
$b_{e,t}^{dis}$	Battery discharging power, e at t
$SOC_{e,t}$	State of charge of battery e at t
$p_{e,t}^{FCRD\uparrow res}$	Up capacity reserved by e at t
$p_{e,t}^{FCRD\downarrow res}$	Down capacity reserved by e at t
$p_{e,t}^{FCRD\uparrow act}$	Up energy actually activated
$p_{e,t}^{FCRD\downarrow act}$	Down energy actually activated
$z_t^{\uparrow}, z_t^{\downarrow}$	Binary bid indicators, up and down

$$R_t^{FCRD\uparrow} = \lambda_t^{FCRD\uparrow} \sum_{e \in E} p_{e,t}^{FCRD\uparrow res} \tag{5}$$

$$R_t^{FCRD\downarrow} = \lambda_t^{FCRD\downarrow} \sum_{e \in E} p_{e,t}^{FCRD\downarrow res} \tag{6}$$

3.3. Per-community power balance

Modified from Crowley et al. (2024) (Eqs. C-1b to C-1d). Because BTMs cannot share power between buildings without incurring the full tax and tariff burden, the power balance is enforced *per community* rather than at the aggregate level. This is one of the main differences from the Crowley et al. formulation. We essentially define net demand $p_{e,t}$ to be used in the battery power balance, assuming non-negative imports $p_{e,t}^{im}$ and exports $p_{e,t}^{ex}$:

$$p_{e,t} - p_{e,t}^{im} + p_{e,t}^{ex} = 0, \qquad \forall e \in E, t \in T \tag{E-1}$$

$$p_{e,t}^{ex} \geq 0, \qquad \forall e \in E, t \in T \tag{E-2}$$

$$p_{e,t}^{im} \geq 0, \qquad \forall e \in E, t \in T \tag{E-3}$$

3.4. Battery power balance

Same structure as Crowley et al. (2024) (Eq. C-1i). Net consumption is covered by PV production ($PV_{e,t}$), demand, i.e. how much energy each energy community uses each hour ($D_{e,t}$), charging/discharging $b_{e,t}^{ch}$, $b_{e,t}^{dis}$, and FCR-D activations $p_{e,t}^{FCRD\downarrow act}$, $p_{e,t}^{FCRD\uparrow act}$. Down activation charges the battery; up activation discharges it:

$$\begin{aligned} &p_{e,t} + PV_{e,t} - D_{e,t} - \left(b_{e,t}^{ch} + p_{e,t}^{FCRD\downarrow act}\right) \\ &+\left(b_{e,t}^{dis} + p_{e,t}^{FCRD\uparrow act}\right) = 0, \quad \forall e \in E, t \in T \end{aligned} \tag{E-4}$$

3.5. FCR-D activation tracking

Replacing the linear interpolation of Crowley et al. (2024) with per 0.1-second fractions since data is available (Fingrid, 2025). Some fraction of each hour the grid frequency f_τ is outside the normal frequency band the TSO wants to keep frequency within, $[49.9, 50.1]$ Hz. If it is above 50.1 Hz, FCR-D down is activated; if it is below 49.9 Hz, FCR-D up is activated. We obtain per 0.1-second frequency data from Fingrid (Fingrid, 2025); for each 0.1-second time step τ within a given hour we calculate the fractions that the frequency deviates in each direction as $y_t^{actFCRD\uparrow}$ and $y_t^{actFCRD\downarrow}$. The activated energy in each direction is then computed as:

$$p_{e,t}^{FCRD\uparrow act} = y_t^{actFCRD\uparrow} p_{e,t}^{FCRD\uparrow res}, \quad \forall e \in E, t \in T \tag{E-5a}$$

$$p_{e,t}^{FCRD\downarrow act} = y_t^{actFCRD\downarrow} p_{e,t}^{FCRD\downarrow res}, \quad \forall e \in E, t \in T \tag{E-5b}$$

3.6. State of charge dynamics

Modified from Crowley et al. (2024) (Eqs. C-2j, C-2k, C-2l). SOC evolves from charging/discharging including the FCR-D activation flows. Battery efficiency is included as η_e. Boundary SOC is fixed at 50% of capacity at the start and end of the horizon:

$$\begin{aligned} SOC_{e,t} = SOC_{e,t-1} + \eta_e\left(b_{e,t}^{ch} + p_{e,t}^{FCRD\downarrow act}\right) \\ - \frac{1}{\eta_e}\left(b_{e,t}^{dis} + p_{e,t}^{FCRD\uparrow act}\right), \\ \forall e \in E,\ t \in T \setminus \{1\} \end{aligned} \tag{E-6a}$$

$$\begin{aligned} SOC_{e,1} = 0.5\,\bar{S}_e + \eta_e\left(b_{e,1}^{ch} + p_{e,1}^{FCRD\downarrow act}\right) \\ - \frac{1}{\eta_e}\left(b_{e,1}^{dis} + p_{e,1}^{FCRD\uparrow act}\right), \forall e \end{aligned} \tag{E-6b}$$

$$SOC_{e,T^{max}} = 0.5\,\bar{S}_e, \quad \forall e \in E \tag{E-6c}$$

Modified from Thingvad et al. (2024) (Eq. T-13). The reserved capacity must be physically deliverable for the required sustain duration. Our portfolio is classified as a Limited Energy Reservoir (LER) asset (Energinet, 2023); the regulatory minimum sustain duration is $\upsilon_E^D = 0.33\,\text{h}$ (20 minutes) (Thingvad et al., 2024). Accounting for battery efficiency η_e: the reserved quantity $p_{e,t}^{FCRD\uparrow res}$, accounting for efficiency and the share we must hold according to LER $\frac{\upsilon_E^D}{\eta_e}$, of the accepted bids $y_t^{accFCRD\uparrow}$, must be less than our state of charge $SOC_{e,t}$. For the up regulation direction:

$$SOC_{e,t} \geq y_t^{accFCRD\uparrow}\,\frac{\upsilon_E^D}{\eta_e}\,p_{e,t}^{FCRD\uparrow res} \tag{E-7a}$$

$$\bar{S}_e - SOC_{e,t} \geq y_t^{accFCRD\downarrow}\,\upsilon_E^D\,\eta_e\,p_{e,t}^{FCRD\downarrow res} \tag{E-7b}$$

Similarly for the down regulation, here not $SOC_{e,t}$ but the headroom to take in more electricity $\bar{S}_e - SOC_{e,t}$.

3.7. Battery power and capacity limits for FCR-D

Directly from Thingvad et al. (2024) (Eqs. T-10, T-11). The battery cannot simultaneously discharge for self-consumption $b_{e,t}^{dis}$ and reserve up regulation $p_{e,t}^{FCRD\uparrow res}$ beyond its rated power $\bar{b}_e$; analogously for charging and down regulation. SOC must remain within $[0, \bar{S}_e]$. Furthermore, LER requirements (Energinet, 2023; Thingvad et al., 2024) mandate that whenever a bid volume is reserved, a fraction $\upsilon_P^D = 0.20$ of that volume $p_{e,t}^{FCRD\downarrow res} \cdot \upsilon_P^D$ must be held available in the opposite direction for NEM to restore SOC after activation. This yields the following inverter power constraints:

$$\bar{b}_e \geq p_{e,t}^{FCRD\uparrow res} + p_{e,t}^{FCRD\downarrow res} \cdot \upsilon_P^D + b_{e,t}^{dis}, \quad \forall e \in E,\, t \in T \tag{E-8a}$$

$$\bar{b}_e \geq p_{e,t}^{FCRD\downarrow res} + p_{e,t}^{FCRD\uparrow res} \cdot \upsilon_P^D + b_{e,t}^{ch}, \quad \forall e \in E,\, t \in T \tag{E-8b}$$

3.8. Minimum bid size

The aggregator either submits aggregated bids $\sum_{e\in E} p_{e,t}^{FCRD\uparrow res}$, $\sum_{e\in E} p_{e,t}^{FCRD\downarrow res}$ within the bid range $[p^{min}, \bar{P}]$ or zero. Binary variables $z_t^{\uparrow}, z_t^{\downarrow} = 0$ allow the bid range to collapse to zero, while $z_t^{\uparrow}, z_t^{\downarrow} = 1$ constrains the bid to be within the given range, yielding a mixed-integer linear program:

$$p^{min}\, z_t^{\uparrow} \leq \sum_{e\in E} p_{e,t}^{FCRD\uparrow res} \leq \bar{P}\, z_t^{\uparrow} \tag{E-10a}$$

$$p^{min}\, z_t^{\downarrow} \leq \sum_{e\in E} p_{e,t}^{FCRD\downarrow res} \leq \bar{P}\, z_t^{\downarrow} \tag{E-10b}$$

$$z_t^{\uparrow},\, z_t^{\downarrow} \in \{0,1\}, \quad \forall t \in T \tag{E-10c}$$

3.9. Non-negativity

Largely an extension from Crowley et al. (2024) (Eqs. C-1m, C-1n). Neither negative energy nor negative bid capacity is accepted. We assume bids, charge, discharge actions, and state of charge cannot be negative:

$$p_{e,t}^{FCRD\uparrow res},\, p_{e,t}^{FCRD\downarrow res},\, b_{e,t}^{ch},\, b_{e,t}^{dis},\, SOC_{e,t} \geq 0 \tag{E-11}$$

4. FCR-D Late Formal Definition

The late auction closes at 18:00 the day before delivery. By that time the early auction has cleared, the day-ahead spot price has been published, and updated forecasts are available. We treat the early and late auctions as separate sequential problems and assume the early accepted quantities are taken as given when solving the late problem. The aggregator may then either *top up* (place additional positive bids) or *buy back* (cancel via negatively oriented bids) the early position. The full derivation of the additional structure relative to Section 3 is in Appendix B and the full listing of equations in Appendix D.

4.1. Additional variable definitions

Some additional parameters are introduced in the late model (Table 3). Largely all parameters are carried over from the early model; we now denote the price variables with a superscript *early* or *late* to differentiate them, and similarly for the accepted volumes. New model parameters are the price to buy back volume in either direction $\{\lambda_t^{FCRD\downarrow bb}, \lambda_t^{FCRD\uparrow bb}\}$, which is the maximum of the late and early auction price for the given instrument.

For the internally solved decision variables we introduce variables to top up reserved volume in the late auction $\{a_{e,t}^{\uparrow}, a_{e,t}^{\downarrow}\}$ as well as buy-back volume $\{c_{e,t}^{\uparrow}, c_{e,t}^{\downarrow}\}$, at the above-stated buy-back prices.

Again we include new binary variables $w_t^{\uparrow}, w_t^{\downarrow}$ in order to account for the non-linearity in either bidding 0 or within the range $[p^{min}, \bar{P}]$.

The final FCR-D up reserved capacity $p_{e,t}^{FCRD\uparrow res}$ used by all physical and battery constraints is given by the early auction accepted bid reservations $p_{e,t}^{FCRD\uparrow acc}$ plus the topped-up volume $a_{e,t}^{\uparrow}$ minus the bought-back volume $c_{e,t}^{\uparrow}$. For

Table 3. Additional sets, parameters and decision variables introduced by the late auction model. All early-auction symbols of Table 2 carry over.

Symbol	Description
Parameters (new)	
$\lambda_t^{FCRD\uparrow early}$	Cleared up price from early auction
$\lambda_t^{FCRD\downarrow early}$	Cleared down price from early auction
$\lambda_t^{FCRD\uparrow late}$	Up price at the late auction
$\lambda_t^{FCRD\downarrow late}$	Down price at the late auction
$\lambda_t^{FCRD\uparrow bb}$	Up buy-back price, $\max\{$early, late$\}$
$\lambda_t^{FCRD\downarrow bb}$	Down buy-back price, $\max\{$early, late$\}$
$p_{e,t}^{FCRD\uparrow acc}$	Up volume accepted in early auction (given)
$p_{e,t}^{FCRD\downarrow acc}$	Down volume accepted in early auction (given)
Decision variables (new)	
$a_{e,t}^{\uparrow}$	Up top-up volume in late auction
$a_{e,t}^{\downarrow}$	Down top-up volume in late auction
$c_{e,t}^{\uparrow}$	Up buy-back (cancellation) volume
$c_{e,t}^{\downarrow}$	Down buy-back (cancellation) volume
$w_t^{\uparrow}, w_t^{\downarrow}$	Binary top-up indicators
$z_t^{\uparrow}, z_t^{\downarrow}$	Binary final-reservation indicators

FCR-D down the same structure applies:

$$p_{e,t}^{FCRD\uparrow res} = p_{e,t}^{FCRD\uparrow acc} + a_{e,t}^{\uparrow} - c_{e,t}^{\uparrow}, \quad \forall e,t \tag{L-1a}$$

$$p_{e,t}^{FCRD\downarrow res} = p_{e,t}^{FCRD\downarrow acc} + a_{e,t}^{\downarrow} - c_{e,t}^{\downarrow}, \quad \forall e,t \tag{L-1b}$$

4.2. Objective Function

Two terms are added relative to the early objective: $C^{FCRD\text{-}bb}$ and $R^{FCRD\text{-}topup}$. Top-ups earn revenue at the late price; buy-backs incur cost at the buy-back price, which is the maximum of the early and late prices $\lambda_t^{FCRD\uparrow bb} = \max\{\lambda_t^{FCRD\uparrow early}, \lambda_t^{FCRD\uparrow late}\}$ (you forgo the early revenue and pay the higher of the two prices to close the position):

$$\begin{aligned}\min \sum_{t\in T} \Big[& C_t^{im} - R_t^{ex} \\ & - R_t^{FCRD\uparrow} - R_t^{FCRD\downarrow} \\ & + C_t^{FCRD\text{-}bb} - R_t^{FCRD\text{-}topup} \Big]\end{aligned} \tag{L-obj}$$

Where the cost of buying back capacity $C^{FCRD\text{-}bb}$ is the buy-back price for up $\lambda_t^{FCRD\uparrow bb}$ and down $\lambda_t^{FCRD\downarrow bb}$ multiplied with how much capacity we are buying back in each direction $\sum_{e\in E} c_{e,t}^{\uparrow}$, $\sum_{e\in E} c_{e,t}^{\downarrow}$:

$$C_t^{FCRD\text{-}bb} = \lambda_t^{FCRD\uparrow bb} \sum_{e\in E} c_{e,t}^{\uparrow} + \lambda_t^{FCRD\downarrow bb} \sum_{e\in E} c_{e,t}^{\downarrow} \tag{7}$$

Similarly the top-up revenue $R^{FCRD\text{-}topup}$ is the late reservation price $\lambda_t^{FCRD\uparrow late}$, $\lambda_t^{FCRD\downarrow late}$ multiplied with the aggregated topped-up volume in each direction $\sum_{e\in E} a_{e,t}^{\uparrow}$, $\sum_{e\in E} a_{e,t}^{\downarrow}$:

$$R_t^{FCRD\text{-}topup} = \lambda_t^{FCRD\uparrow late} \sum_{e\in E} a_{e,t}^{\uparrow} + \lambda_t^{FCRD\downarrow late} \sum_{e\in E} a_{e,t}^{\downarrow} \tag{8}$$

4.3. Minimum bid size for top-ups

Any *new* positive volume submitted to the late auction must satisfy the minimum bid constraint. This produces a similar non-linearity as seen in Section 3, where the binary variables here are $w_t^{\uparrow}$, $w_t^{\downarrow}$:

$$p^{min}\, w_t^{\uparrow} \le \sum_{e\in E} a_{e,t}^{\uparrow} \le \bar{P}\, w_t^{\uparrow} \tag{L-2a}$$

$$p^{min}\, w_t^{\downarrow} \le \sum_{e\in E} a_{e,t}^{\downarrow} \le \bar{P}\, w_t^{\downarrow} \tag{L-2b}$$

$$w_t^{\uparrow},\, w_t^{\downarrow} \in \{0,1\} \tag{L-2c}$$

4.4. Minimum bid size for final reservation

The total final reservation per direction must also either be zero or above p^{min}. Without this, partial cancellation could leave a stranded position below the minimum bid floor:

$$p^{min}\, z_t^{\uparrow} \le \sum_{e\in E} p_{e,t}^{FCRD\uparrow res} \le \bar{P}\, z_t^{\uparrow} \tag{L-3a}$$

$$p^{min}\, z_t^{\downarrow} \le \sum_{e\in E} p_{e,t}^{FCRD\downarrow res} \le \bar{P}\, z_t^{\downarrow} \tag{L-3b}$$

$$z_t^{\uparrow},\, z_t^{\downarrow} \in \{0,1\} \tag{L-3c}$$

4.5. No simultaneous top-up and buy-back

For a given direction and hour, the aggregator cannot both add new volume and cancel existing volume; if $w_t = 1$ (we top up) then the maximum buy-back collapses to zero. Otherwise the maximum buy-back is capped by what was accepted in the early auction:

$$c_{e,t}^{\uparrow} \le p_{e,t}^{FCRD\uparrow acc}\,(1 - w_t^{\uparrow}), \quad \forall e,t \tag{L-4a}$$

$$c_{e,t}^{\downarrow} \le p_{e,t}^{FCRD\downarrow acc}\,(1 - w_t^{\downarrow}), \quad \forall e,t \tag{L-4b}$$

4.6. Carried-over constraints from the early model

All physical constraints from Section 3 continue to hold, with $p_{e,t}^{FCRD\uparrow res}$ and $p_{e,t}^{FCRD\downarrow res}$ now defined by the linking equations (L-1a) and (L-1b):

- Per-community power balance (E-1), (E-2), (E-3).

- Battery power balance (E-4).
- Activation tracking (E-5a), (E-5b).
- State of charge dynamics (E-6a)–(E-6c).
- Sustain feasibility (E-7a), (E-7b).
- Battery power and capacity limits (E-8a), (E-8b).
- Non-negativity of $p_{e,t}^{FCRD\uparrow res}$, $p_{e,t}^{FCRD\downarrow res}$, $b_{e,t}^{ch}$, $b_{e,t}^{dis}$, $SOC_{e,t}$, $a_{e,t}^{\uparrow}$, $a_{e,t}^{\downarrow}$, $c_{e,t}^{\uparrow}$, $c_{e,t}^{\downarrow}$.

5. Case Study: Simulated 1 MW Portfolio

5.1. Portfolio construction

We base our simulated portfolio on empirical energy community data, here two sets, s and b, which each contain a consumption and production profile shape. Each simulated EC is a scaled version of one of these curves with some spread η_σ. We assume the EC profile scale factors are $X \sim \mathcal{N}(\mu = 1, \sigma = 1)$. We draw type b with probability P_b and type s with probability P_s, exclude negative outcomes, and then apply these scale factors to obtain simulated EC production and consumption curves.

Table 4. EC portfolio configuration

Parameter	Value	Unit
Number of ECs	10	–
Max power per EC	100	kW
Capacity per EC	200	kWh
Round-trip efficiency η	0.95	–
Initial / terminal SOC	0.50	fraction
Total portfolio power	1000	kW
Total portfolio capacity	2000	kWh
Min. FCR-D bid size $P_{\min}$	100	kW
Sustain duration υ_E^D	1/3	h
Prob. b-type P_b	0.60	–
Prob. s-type P_s	0.40	–
Scale factor mean μ	1.0	–
Scale factor std σ	1.0	–

5.2. One day run results: May 11th, 2025

To get some intuition for what choices the algorithm makes we look at one day here. For our synthetic portfolio we can see the variation in consumption and production as well as the mean for each type in Figure 6. This is the basis on which the optimisation is made. Each of these 10 energy communities has a battery as specified in Table 4, that together add up to 1 MW aggregated power.

Furthermore we solve the linear program for the example day to illustrate the choices of when to charge, discharge and reserve capacity that we will make to minimise the objective given the constraints. Arbitrage-related information for the optimal solution is illustrated in Figure 6. Solar production peaks during the day hours 06:00–17:00, and consumption is lower during the night and higher during the day. This figure, although it does not display it explicitly, also takes FCR-D up and down bidding into account, which also affects the optimal SOC for the battery. However, we still see clearly that the traditional strategy of charging during the day where we have excess solar and relatively lower prices, and then discharging during the afternoon where prices are the highest, to cover consumption, still is a dominant strategy and supersedes the alternative of keeping SOC more level in order to reserve more capacity in the ancillary services market.

In Figure 6 we observe the adjustment that takes place in the late auction program, where we adjust the accepted bids from the morning either up or down for each hour. The program essentially takes the same input predictions as the morning program, with the addition of taking the early accepted quantities as given and then minimising the objective now including the cost of buying back (i.e. adjusting down a given bid for an hour) and the revenue for topping up. Figure 6 also shows the reserved quantities for FCR-D up and FCR-D down each hour after each auction. We imagine the first auction is run on a noisy forecast while the late auction is somewhat more reliable, also since the real day-ahead spot price will have been published before it is submitted, giving some more informational certainty relative to the early forecasts.

Table 5. Single-day summary — 2025-05-11

Stream	DKK	Share
No-battery cost	2182.4	—
Base cost (battery only)	416.5	—
Battery arbitrage saving	1766.0	51.9%
Arb improvement from FCR-D	-182.2	-5.4%
FCR-D Up (early+late-bb)	1167.2	34.3%
early reservation	1218.8	—
late top-up	27.4	—
buy-back cost	-79.1	—
FCR-D Down (early+late-bb)	650.1	19.1%
early reservation	386.9	—
late top-up	299.3	—
buy-back cost	-36.1	—
TOTAL VALUE	**3401.1**	**100.0%**

The choices made in the optimal solutions appear rational. We see for an example that for FCR-D down the late price is much higher than the early price, on which the initial bids are based, in the time span 08:00–15:00. This causes the late program to top up down during 10:00–15:00. We must

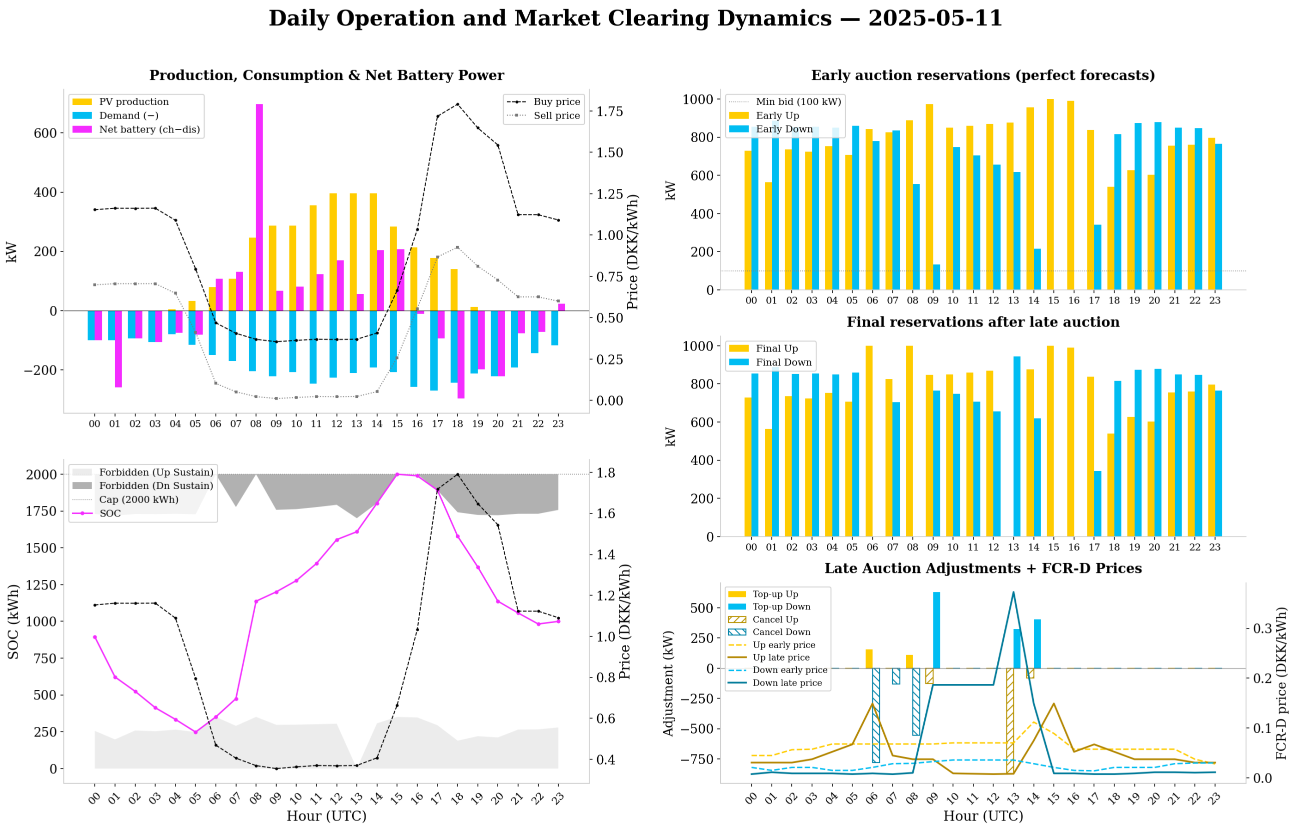

Figure 6. Optimal charge/discharge schedule and SOC trajectory together with spot prices for May 11th, 2025. As well as the relationship between SOC and the reserved capacity in the aggregated battery at each hour, illustrating the trade-off between ancillary services participation and self consumption / arbitrage. The only information update between early and late with perfect forecasts is the FCR-D late price. We observe that we adjust up when the late FCR-D down price is high 12:00–14:00.

still remember here that we also take self consumption and arbitrage into account, and if we look at the SOC chart in Figure 6 we observe that the SOC hits max at 15:00 where the price also starts to drop again, causing top-up behaviour to stop and self consumption battery discharge to start.

Figure 6 illustrates how SOC relates to the reserved capacity in the aggregated battery at each hour. Here we can clearly see the fundamental trade-off between reserving capacity for ancillary services and having more range for SOC to contribute to self consumption and arbitrage. Given that we have both up and down FCR-D, a good strategy for the battery is to remain close to half SOC, where it can maximise reservation capacity in both directions. Think of it like this: when the battery is exactly in the middle we get equal free space for up and down, and hence we have two revenue streams from FCR-D; however as soon as we start to deviate from 50% in any direction we improve capacity for one direction at the expense of the other. The additional revenue we gain from this deviation is then only one revenue stream, not two. Hence for it to be rational this additional revenue will have to also cover the loss of revenue from participation in the other direction, making the rational behaviour most of the time to stay near 50% SOC.

It is then plausible that we see a more stable SOC during most hours and then a sharper charge to go to 100% before the electricity price rises, to cover self consumption during the expensive evening spot prices. This signifies a shift in strategy priority from ancillary services to self consumption during these critical hours, because the high spot price makes it more profitable to self consume than to maximise reserved capacity during the hours 15:00–21:00.

5.3. One year run results: 2025

We run the daily program for each day in 2025 given the fictional 1 MW aggregated portfolio.

Figure 7 breaks down the cost spent on energy and the adjustment volume from the late auction. It also breaks down the revenue between the early and late auctions. We

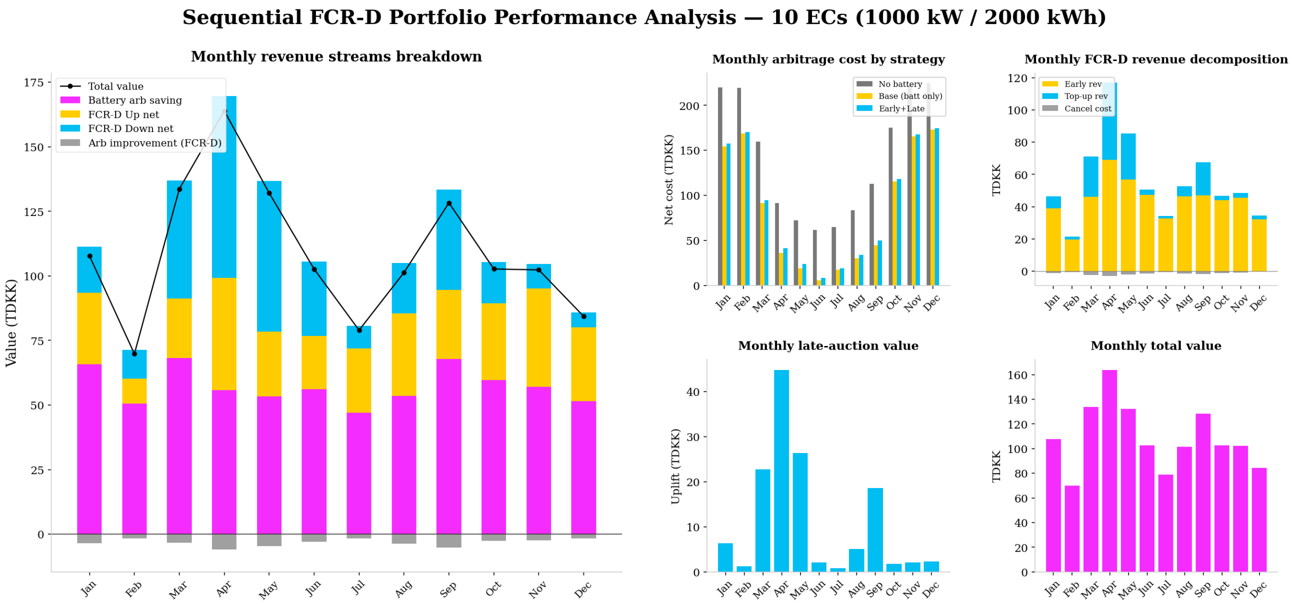

Figure 7. Monthly breakdown of the 2025 backtest. In the left figure we observe arbitrage savings in magenta and how they degrade when we stack FCR-D on top as the grey negative mass under each staple. We observe substantial FCR-D revenue in all months. Right figure shows again the cost reduction of having a battery and then a slight increase for FCR-D participation; it also seems that late revenue is relatively volatile in comparison to early.

observe that having the battery lowers the cost spent on electricity significantly, and stacking early and late auction participation only slightly alters the cost, meaning that FCR-D market participation does not significantly hurt the regular self consumption and arbitrage benefits that exist. Another observation is that the cancellation cost we spend in the second auction remains low, pointing towards a strategy where we usually let bids from the early auction stand, or top them up. This is logical since cancellation often is unprofitable because the buy-back price is the maximum of the early and late auction prices. Generally the early auction provides the bulk of the revenue, but the late auction participation is still very worthwhile and has a much higher variance compared to the early auction revenue.

Figure 7 and table 7 also shows the different revenue streams over the months of the year, comparing arbitrage and self consumption to the case where we have no battery at all, and the stacked revenue for FCR-D up and down. Interestingly, we observe that the improvement in arbitrage when stacking ancillary services is negative, but again only slightly in comparison with the ancillary services revenue which is many times greater for each month. With this we can conclude that it is possible to make a profit with BTM-type communities participating in these markets based on the 2025 data. This of course assumes the fact that this is an upper bound based on perfect forecasts, which will not be the case for real market participation. Nonetheless it signals

Table 6. Yearly Summary 2025 — 365 days (10 ECs, 1000 kW / 2000 kWh)

Stream	**TDKK**	**Share**
Battery arbitrage saving	686.4	52.5%
Arb improvement from FCR-D	-38.9	-3.0%
FCR-D Up net	329.6	25.2%
early reservation	301.5	—
late top-up	36.3	—
buy-back cost	-8.2	—
FCR-D Down net	330.8	25.3%
early reservation	224.5	—
late top-up	114.8	—
buy-back cost	-8.4	—
TOTAL VALUE	**1307.9**	**100.0%**

Table 7. Monthly Revenue Breakdown (TDKK)

Month	Arb	ΔArb	Up	Down	Total
Jan	65.9	-3.5	27.5	17.9	107.8
Feb	50.5	-1.6	9.7	11.3	69.8
Mar	68.1	-3.2	23.2	45.6	133.7
Apr	55.7	-5.8	43.6	70.4	163.9
May	53.4	-4.7	24.9	58.4	132.1
Jun	56.2	-2.9	20.5	28.8	102.6
Jul	47.1	-1.7	24.8	8.8	79.0
Aug	53.6	-3.7	31.9	19.5	101.3
Sep	67.8	-5.2	26.8	39.0	128.3
Oct	59.6	-2.6	29.8	15.9	102.7
Nov	57.0	-2.3	38.2	9.5	102.3
Dec	51.5	-1.6	28.7	5.8	84.3

that there is a theoretical incentive for the BTM asset type to participate in ancillary services markets.

Furthermore Table 6 shows that the total value from the battery based on the 2025 backtest with perfect forecasts was about 1.3 million DKK, where the stacked ancillary services FCR-D up and down stood for 50.5% of the earnings. This shows that ancillary services participation has the ability to significantly contribute to the attractiveness of installing batteries in BTM-style energy communities and that it can be done without significantly harming self consumption savings, as the loss of self consumption arbitrage earnings was only 38.9 TDKK.

6. Portfolio Variations

In order to observe how results vary depending on portfolio construction see results in figure 8, six different variations were run see configuration in table 8 and detailed results in table 9. C1 is meant to test whether it is viable to bid into FCR-D with a portfolio that just clears the minimum requirement. C2 is a large heterogeneous portfolio of varying sizes meant to simulate a realistic aggregator scenario. C3 is the baseline from the previous one-year study. C4 mirrors the C3 baseline but with half the battery size. C5 simulates fewer but larger energy communities, and C6 tests how performance degrades when end-of-life batteries with lower efficiency are included.

Across the board the FCR-D revenue share remains large, ranging from 36 to 56 percent, showing the general attractiveness of FCR-D participation for BTM portfolios. Largely it seems that portfolios can expect to earn around 600 DKK per kWh and around 1,250 DKK per kW for similar portfolio setups. The smaller portfolio also achieved higher returns per kW and per kWh than average, capturing only the top opportunities in one direction at a time. The fact that reducing the power or battery size resulted in a higher return per kW for C1 and per kWh for C4 shows that there may be some diminishing returns to FCR-D participation, where a more constrained portfolio can choose to participate only in the best opportunities, and there is some marginality to the size of all opportunities presented to an unconstrained portfolio.

The reduction in efficiency for C6 only produced a small performance degradation compared to C3. The efficiency drop requires more kW and kWh in order to place and fulfil bids. Since the decline only affected some batteries, it is natural that the overall effect was relatively small.

The baseline community C3 beats the large modern C5 by a slight margin, mostly due to higher arbitrage revenue in C3, as the FCR-D share was slightly higher for C5 at 56 percent compared to 51 percent. The main reason is likely that the PV and load are significantly smaller in C5, meaning there is less consumption to offset.

Table 8. Portfolio configuration parameters and annual energy profiles. Load and PV are annual totals for the simulated portfolio.

Config	ECs	kW	kWh	η	Load (MWh)	PV (MWh)
Barely viable	10	100	200	0.95	1,618	600
Heterogeneous	20	1534	3514	0.95	3,065	1,154
Baseline	10	1000	2000	0.95	1,618	600
Thin battery	10	1000	1000	0.95	1,618	600
Large modern	5	1000	3000	0.95	741	301
Mixed η	10	1000	2000	0.80–0.95	1,618	600

Note: The minimum bid limit ($P_{\min}$) is a uniform 100 kW across all configuration profiles.

7. Revenue Forecast Noise Sensitivity

After deriving an upper bound through perfect forecast assumptions we now challenge this by introducing noise to the forecasts. The methodology is inspired by Langtry et al. (2024), which examines forecast sensitivity of a linear program to different levels of noise across individual variables in a model predictive control (MPC) context for building energy management. The results can be seen in figure 9.

7.1. The Noise Model

We test sensitivity by adding a Gaussian random walk (GRW) noise component to ground-truth values.:

$$f^{v}_{\text{GRW}}[t,\tau] = v_{t+\tau} + \sum_{j=1}^{\tau} w_j, \quad w_j \sim \mathcal{N}(0,\sigma^2), \quad (9)$$

where $\sigma = \alpha\hat{\sigma}$ and $\hat{\sigma}$ is the empirical average hourly standard deviation of a persistence forecast error:

$$\hat{\sigma} = \text{mean}_h \big(\text{std}(v[d+1,h] - v[d,h])_{d\in 2025}\big)\,. \quad (10)$$

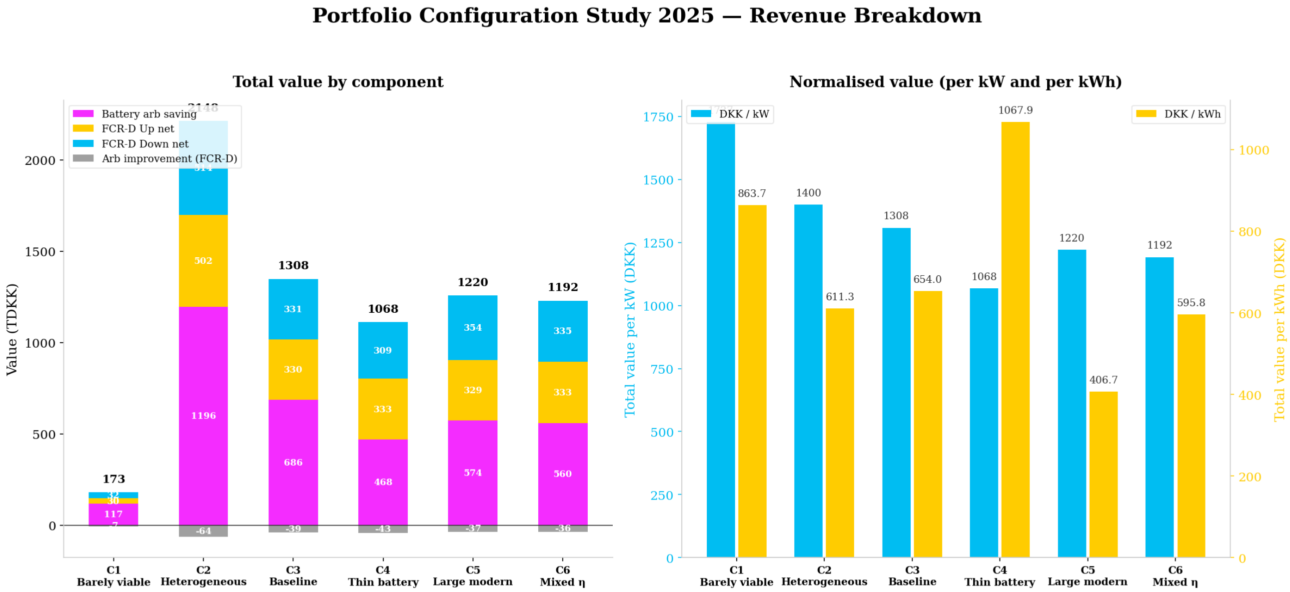

Figure 8. Left figure shows the total revenue achieved on the 2025 back test broken down into its components. The base C3 case is equivalent to our earlier test. The right figure compares how much the different portfolio configurations earn per kW and kWh based on the same back test.

Table 9. Portfolio Study 2025 — Annual Summary Breakdown

Configuration	Power (kW)	Cap. (kWh)	Arb. (TDKK)	ΔArb. (TDKK)	FCR-D↑ (TDKK)	FCR-D↓ (TDKK)	Late Rev. (TDKK)	Buyback (TDKK)	Total (TDKK)	FCR-D Share (%)	per kW (DKK)	per kWh (DKK)
Barely viable	100	200	117.3	−7.4	30.4	32.4	25.0	−3.0	172.7	36.4	1727.3	863.7
Heterogeneous	1534	3514	1195.7	−63.6	502.2	513.8	250.1	−26.0	2148.1	47.3	1400.3	611.3
Baseline	1000	2000	686.4	−38.9	329.5	330.9	151.1	−16.6	1307.9	50.5	1307.9	654.0
Thin battery	1000	1000	468.0	−42.9	333.4	309.4	135.0	−15.4	1067.9	60.2	1067.9	1067.9
Large modern	1000	3000	573.6	−37.2	329.2	354.4	168.1	−18.1	1220.0	56.0	1220.0	406.7
Mixed efficiency	1000	2000	560.0	−36.4	333.4	334.6	149.4	−16.5	1191.7	56.1	1191.7	595.8

The cumulative structure ensures forecasts are smooth and progressively noisier for longer horizons. The auto-correlated noise model is not intended to accurately simulate a specific ML forecast, but rather to give a comparative analysis of how sensitive revenue is to noise in different variables.

7.2. Portfolio Construction with Noise

Consumption noise is applied individually per building, while PV forecast noise depends primarily on weather reports and is applied identically across buildings (but to each individual production curve). All other variables (prices and activation fractions) are fed in at the aggregate level. We also include two simple persistence forecasts a 2 day and 7 day persistence, meaning we forecast a given hour as the value at the same hour 2 or 7 days ago. Figure 4 shows the cyclical nature of many of our variables and why persistence may work reasonably good particularly in PV and load forecasting. The revenue results in percent of the perfect forecast revenue is also plotted alongside the noised variables to compare a real although simple and rule based forecast method to the artificial noise and establish a lower bound in terms of revenue expectations compared to more advanced forecasting methods.

7.3. Noise Sensitivity Results

Observing the results of the noise study in Figure 9, we conclude that persistence forecasts hold up well, with D-2 persistence producing 86.6 percent of the perfect-forecast revenue and D-7 persistence producing 88.2 percent.

We also observe that introducing noise into the load variable has the most negative impact on revenue, while the other variables do not degrade performance as much. This is possibly because the load variable is derived individually per energy community, making forecast errors harder to diversify away at the portfolio level.

Most surprisingly, we observe that introducing noise into the

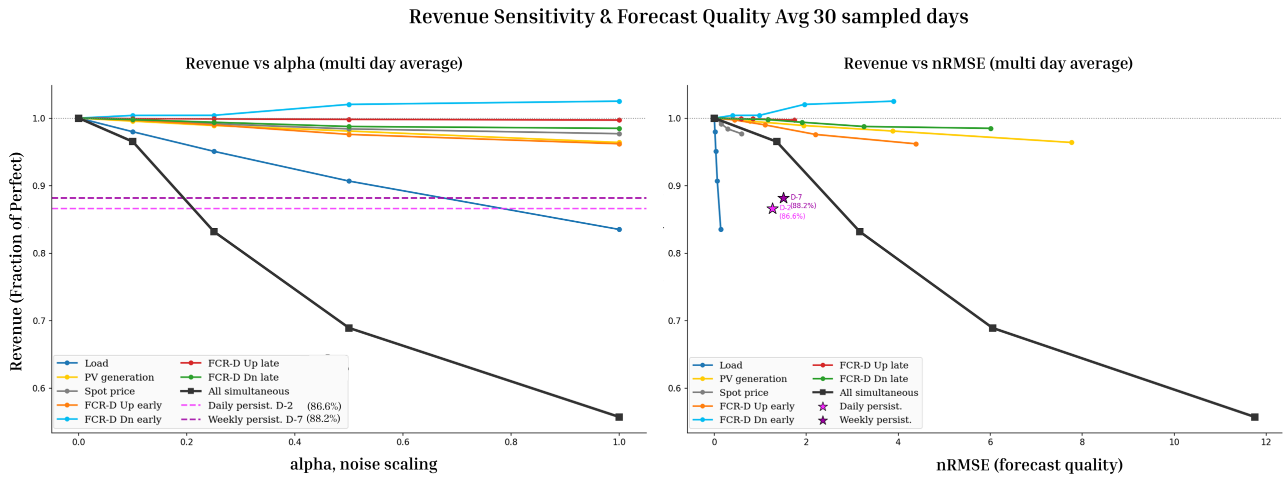

Figure 9. Noise is scaled through alpha that increases the standard deviation of the noise added each step. It appears that noise in the consumption forecast degrades performance the most compared to other variables. The effect is also larger where all variables are noisy.

FCR-D down early price produces higher revenue than the perfect-forecast case, something that should be impossible given a correct model specification. The reason this can occur is that the model does not constitute a mathematical upper bound, since it is executed in two sequential stages. A scenario that can produce this outcome is one where the FCR-D up early price is very high and the FCR-D down early price is very low, causing the model to adjust the state of charge to make room for more FCR-D up reservation in the early auction. If the FCR-D down price then becomes attractive in the late auction, the model is unable to reserve the optimal quantity of down capacity because the SOC has already been committed to up. When noise is introduced into the FCR-D down early prices, which are often near zero, the clip at zero creates a positive bias that makes these prices appear more attractive than they truly are. This reduces the tendency to over-commit SOC toward up in the early auction, leaving more room for down reservation in the late auction. While this reveals a suboptimality in the sequential structure, it is plausible that simple heuristic rules could produce a similar corrective effect in practice.

$$\text{nRMSE} = \frac{\frac{1}{N}\sum_{t=0}^{N-1}\sqrt{\frac{1}{T}\sum_{\tau=1}^{T}(f_{t,\tau}^{v} - v_{t+\tau})^2}}{\frac{1}{N}\sum_{t=0}^{N-1} v_t} \tag{11}$$

where N is the number of time instances, T is the forecasting horizon length, $f_{t,\tau}^{v}$ is the forecast of variable v at time t for horizon step τ, $v_{t+\tau}$ is the corresponding true value, and $\frac{1}{N}\sum_{t=0}^{N-1} v_t$ is the mean level of the target variable.

8. Discussion

The 2025 backtest with perfect forecasts seems to stack well with the pre-existing goals of self-consumption and arbitrage for the aggregated BTM community asset. We observe this in the fact that the self-consumption revenue is not significantly hurt by imposing this additional goal, the revenue loss was around 38.9 TDKK or about 6 percent. In addition we see a significant revenue increase with FCR-D participation. Under perfect-forecast assumptions, the total portfolio value reached 1,308 TDKK, of which FCR-D up contributed 25.2% and FCR-D down contributed 25.3%, together constituting a 48% uplift over the battery-only baseline. This makes revenue stacking in this scenario particularly attractive because we do not significantly hurt the base economic case upon which we made the initial battery investment, and bring significant upside. The fact that batteries in BTM communities are already installed in order to increase self-consumption also makes this particular setup low risk, while also being suitable from a societal point of view, since these assets now can serve multiple purposes.

The portfolio variation study (Figure 8) showed firstly that significant FCR-D revenue can be earned with different portfolio constellations. We can hypothesise that the market conditions provide some set of opportunities that may be marginal in character, explaining why the more restrained portfolios, while having a somewhat lower revenue, still have a relatively larger FCR-D revenue than expected, possibly by cherry-picking these high-value opportunities. Overall the biggest takeaway is that the revenue stacking seemed relatively stable across portfolios and does not significantly diminish in any constellation we have seen. Even for the minimum bid size portfolio, FCR-D stacking provided around a 35 percent revenue share.

Studying the programme's schedule for the 11th of May case study (Figure 6), we can interpret the schedule as roughly enacting two strategies. One being to keep at approximately 50 percent SOC where large arbitrage opportunities do not warrant other behaviour. The reason for this is that since we can bid up and down FCR-D at the same time in a fixed-capacity system, it means that in order to profitably deviate, the loss of reservation payments in one direction must not only be exceeded by but made up for by the increase in reservation payments from the reservation in the opposite direction, usually requiring the price to be more than double in one direction, usually leading to the 50 percent strategy where we can bid into both markets equally. The second strategy can be observed in the evening where both consumption and prices rise at the same time. This provides a great opportunity for arbitrage, filling the battery up before the peak and then slowly discharging during the evening. We see this goal taking over priority, pushing FCR-D actions into a less optimal state during these hours to allow this arbitrage SOC movement.

The noise study (Figure 9) revealed that revenue figures around 80–90 percent of those achieved with the perfect forecasts are achievable using simple persistence forecasts for two or seven days. This may be one of the main findings of this study and is greatly encouraging for smaller, less sophisticated actors in this space, as persistence often is used as a lower bound for what more advanced machine learning algorithms can achieve. Furthermore, this may point towards ordinal structure amongst the variables mattering more than cardinal value changes. For example, it may be more important which of two prices is higher than how much higher a given price is, making the model less sensitive to noise as long as the values approximately reflect these relationships accurately, which a persistence forecast would do. In addition we find real evidence for the suboptimality in the two-stage problem definition, where the fact that we do not include the late price forecast in the early solver and vice versa leads to possibly over- or under-committing in the early auction, changing SOC in a way that suboptimally restricts late bidding. This could be mitigated somewhat by adding positive noise, mitigating the effect of one FCR-D early price being much higher than the other, causing these drastic SOC movements. Realistically it could also be the case that a rule-based mechanism could add a similar protection against these drastic actions in the early auction if these are a known weak point of the model. Overall revenue was significantly more sensitive to noise within the load forecast, possibly since this is applied per building whereas the other forecasts are mainly applied in aggregate, where solar PV is an in-between case where the curves are building-specific but the noise itself is aggregate-derived to reflect a collective solar forecast for DK2. Not surprisingly we also see that the model is significantly more robust to adding noise to one variable than to having noise added to all variables at once. While this may be seen as a significant limitation, we can also contextualise this with the fact that FCR-D prices are notoriously volatile and difficult to forecast, putting into question how much better the model would perform in practice if it allowed forecasts of late FCR-D prices in the early auction solver as well and vice versa. Furthermore, most literature today, see (Thingvad et al., 2024), simply assumes all volume is traded in the early market, which is a greater simplification; (Crowley et al., 2024) also only simulate one FCR-D auction step.

One of the model's strongest assumptions is full bid acceptance, motivated by the fact that we bid at price zero, since the marginal cost of participation is negligible: battery degradation from FCR-D activations is minimal given that actual activated volumes in practice are very small, and since these assets are already installed and financially motivated by self-consumption and arbitrage objectives. One could hypothesise that oversaturation of battery energy storage systems (BESS) participating in this market could force partial bid acceptance even at price zero. However, this can be contextualised with 2025 data: Energianalyse & PlanEnergi (2025) estimates a DK2 procurement need of 42 MW for both FCR-D up and FCR-D down, while Nordic TSOs (2025) report prequalified battery volumes of only 18.1 MW for each direction as of January 1, 2025. This indicates an undersupply of battery capacity relative to procurement demand, suggesting that zero-price bids from a 1 MW portfolio would face minimal competition and near-certain acceptance during the study period. However, it is possible that the market becomes more saturated in the future, making the full bid acceptance assumption more problematic. The reason this was left out is the lack of public data, making it impossible to model.

Battery degradation is also omitted from this study. We only backtest on one year and this is assumed to be marginal for this time period. However, we include an efficiency variable through which degradation could possibly be fed in if expanding the model further. In the portfolio setup study we also include a test where we include some more degraded batteries, seeing only a marginal decline in revenue compared to the base case.

In 2025 Nord Pool switched to 15-minute resolution for its spot prices, which means that our hourly model acts suboptimally in this regard. If we cannot convert all variables to per 15 minutes, one hybrid approach may be to execute the model based on the highest hourly price and then greedily choose to bid in the 15-minute interval with the highest price within each hour, or to execute similar two-stage strategies.

Note that reference literature (Thingvad et al., 2024) and (Crowley et al., 2024) also use per-hour data, as fully moving to per 15 minutes would require interpolation in all other variables.

Finally, this is a narrow study. The regulations for FCR-D are particular to DK2 and contextualised within the Nordic grid context. While BTM-like setups can be seen in Germany and Sweden, the specificity of this study may make it difficult to extrapolate from.

Regarding future work, one interesting direction would be to extend the model to include both early and late FCR-D prices to provide an accurate upper bound for revenue, as well as contextualising this with the practicality and feasibility of executing this type of programme in practice. Another possible direction would be to extend the programme to include BRP-status-related instruments such as FCR-N, aFRR, and mFRR, to compare the trade-off for an aggregator pursuing one status over the other. A third direction could be to build on this programme as a foundation, adding complementary strategies for hedging risks and constraining behaviour in a practically desirable and safe way so that guarantees can be made.

To conclude, this study has shown good feasibility for aggregated BTM-type assets with batteries to participate in and stack FCR-D revenue on top of existing self-consumption and arbitrage objectives, providing significant upside with very limited harm across multiple portfolio variations. Furthermore, the strong performance of a simple persistence forecast yielding 80–90 percent of perfect-forecast revenue signals that this revenue stream is accessible even with simpler forecasting methods. Although the sequential two-stage model does not constitute a perfect mathematical upper bound, it provides a practical baseline framework for sequentially bidding into ancillary services markets, extending the current Danish FCR-D literature by explicitly modelling both auction stages.

References

Amini Toosi, H., Lavagna, M., Leonforte, F., Del Pero, C., and Aste, N. Building decarbonization: Assessing the potential of building-integrated photovoltaics and thermal energy storage systems. *Energy Reports*, 8:574–581, 2022. doi: 10.1016/j.egyr.2022.10.322.

Committee on Climate Change. The sixth carbon budget: The UK's path to net zero. Technical report, Committee on Climate Change, December 2020. URL https://www.theccc.org.uk/wp-content/uploads/2020/12/The-Sixth-Carbon-Budget-The-UKs-path-to-Net-Zero.pdf.

Crowley, B., Magid, S. E. C., Nielsen, D. B. V., Shin, J., and Skak-Iversen, A. S. Can energy communities provide ancillary services? A stochastic program for aggregated bidding in the Nordic ancillary markets. In *Proceedings of the 2024 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*, pp. 283–288. IEEE, 2024. doi: 10.1109/SmartGridComm60555.2024.10738093.

Energianalyse, E. and PlanEnergi. Økonomiske fordele ved elektrificering og samspil med elmarkedet. Technical report, Energi på Tværs / Region Hovedstaden, October 2025. URL https://gate21.dk/wp-content/uploads/2025/11/Elektrificering-i-hovedstadsregionen-oekonomiske-pdf. Data reference: Energinet. Accessed: June 2026.

Energinet. Prequalification of FCR-N and FCR-D. https://en.energinet.dk/media/r4rbirq2/prequalfication-of-fcr-n-and-fcr-d.pdf, May 2023. Version 1.0, accessed: June 2026.

Energinet. FCR-N and FCR-D market data DK2 (FcrNdDK2). https://www.energidataservice.dk/tso-electricity/FcrNdDK2, 2025a. Accessed: 2026-06-03. Hourly price data for FCR-D Up and FCR-D Down markets.

Energinet. Day-ahead prices (DayAheadPrices) – 15-minute resolution. https://energidataservice.dk/tso-electricity/DayAheadPrices, 2025b. Accessed: 2026-06-03. Used for the second part of 2025 following the structural transition to 15-minute resolution pricing.

Energinet. Elspot prices (Elspotprices) – hourly resolution. https://www.energidataservice.dk/tso-electricity/Elspotprices, 2025c. Accessed: 2026-06-03. Used for the first part of 2025 prior to the 15-minute market time unit transition.

Energinet, Fingrid, Statnett, and Svenska kraftnät. Nordic FCR balancing capacity market – implementation guide for BSP. Technical report, Nordic TSOs, February 2026. https://en.energinet.dk/media/chzh45wk/nordic-fcr-balancing-capacity-market-implementati pdf.

Enyday ApS. Case reference: Vendsysselhus – Fælles solpaneler, batterier og EV-opladere hjælper lejerne i Vendsysselhus mod et mere bæredygtigt forbrug. https://enyday.com/vendsysselhus-dk/, March 2024. Accessed: June 2026.

Enyday A/S. Behind-the-meter energy community consumption and generation profiles, 2025. Private dataset provided in collaboration with Enyday. Contains per-meter energy consumption and solar PV production data for two Danish energy communities.

European Union. Directive (EU) 2018/2001 of the European parliament and of the council of 11 December 2018 on the promotion of the use of energy from renewable sources. *Official Journal of the European Union*, 328: 82–209, 2018.

European Union. Directive (EU) 2019/944 of the European parliament and of the council of 5 June 2019 on common rules for the internal market for electricity and amending directive 2012/27/EU. *Official Journal of the European Union*, 158:125–199, 2019.

Fingrid. Frequency – historical data (dataset id: 339). https://data.fingrid.fi/en/datasets/339, 2025. Accessed: 2026-06-03. Hourly activation fractions derived from 0.1-second resolution frequency data for the Nordic synchronous area (DK2).

International Energy Agency. Tracking clean energy progress 2023. Technical report, IEA, Paris, July 2023. URL https://www.iea.org/reports/tracking-clean-energy-progress-2023.

Langtry, M., Wichitwechkarn, V., Ward, R., Zhuang, C., Kreitmair, M. J., Makasis, N., Xuereb Conti, Z., and Choudhary, R. Impact of data for forecasting on performance of model predictive control in buildings with smart energy storage. *arXiv preprint arXiv:2402.12539*, 2024.

Nordic TSOs. Batteries in the nordic reserve markets. Technical report, Energinet, Fingrid, Statnett, Svenska kraftnät, April 2025. URL https://www.statnett.no/globalassets/for-aktorer-i-kraftsystemet/marked/reservemarkeder/nordic-tsos---batteries-in-the-nordic-reserve-markets-april-2025.pdf. Accessed: June 2026.

Radius Elnet A/S. Netabonnement og tariffer for elkunder. https://radiuselnet.dk/priser/, 2025. Accessed: 2026-06-03. Time-of-Use (ToU) network distribution tariffs used to calculate hourly network import and grid utilization costs for the BTM community.

REScoop.eu. REC & CEC definitions: Transposition tracker. https://www.rescoop.eu/policy/transposition-tracker/rec-cec-definitions, April 2024. Accessed: June 2026.

Thingvad, A., Hameed, Z., Ziras, C., and Noureddine, A. F. Optimal scheduling of limited energy reservoir technologies in Nordic ancillary services and energy markets. In *Proceedings of 2024 International Conference on Renewable Energies and Smart Technologies (REST)*. IEEE, 2024. doi: 10.1109/REST59987.2024.10645417.

ur Rehman, U., Faria, P., Gomes, L., and Vale, Z. Future of energy management models in smart homes: A systematic literature review of research trends, gaps, and future directions. *Process Integration and Optimization for Sustainability*, 9(4):1169–1198, 2025. doi: 10.1007/s41660-025-00506-x.

Williams, H. P. *Model Building in Mathematical Programming*. Wiley, 5th edition, 2013.

Appendix

A. Derivation of the Early FCR-D Model from Crowley et al.

This appendix shows step-by-step how the early model in Section 3 is obtained from the FCR-D base case of Crowley et al. (2024). We start from Crowley's self-consumption base case, layer on FCR-D up, and then introduce the modifications required for the BTM aggregation setting and the addition of FCR-D down. Finally, we incorporate the Limited Energy Reservoir (LER) constraints mandated by the Energinet prequalification rules (Energinet, 2023), following the formulation of Thingvad et al. (2024).

A.1. Crowley's self-consumption base case

Crowley's base program (Eqs. 1a–1o in the original) minimises expected import cost minus export revenue across stochastic scenarios $s \in S$ with probabilities π_s:

$$\min \sum_{s\in S} \pi_s \left[\sum_{t\in T} \lambda^{im}_{t,s}\, p^{im}_{t,s} - \lambda^{ex}_{t,s}\, p^{ex}_{t,s} \right] \tag{C-1a}$$

subject to an aggregate power balance over community members $c \in C$:

$$\sum_{c\in C} p_{c,t,s} - p^{im}_{t,s} + p^{ex}_{t,s} = 0, \quad \forall t \in T \tag{C-1b}$$

with $p^{im}_{t,s}, p^{ex}_{t,s} \geq 0$ (C-1c, C-1d). For battery-prosumer members $e \in E$ the per-member power balance is

$$p_{e,t,s} + PV_{e,t,s} - D_{e,t,s} - b^{ch}_{e,t,s} + b^{dis}_{e,t,s} = 0 \tag{C-1i}$$

and the SOC dynamics are

$$SOC_{e,t,s} = SOC_{e,t-1,s} + \eta\, b^{ch}_{e,t,s} - \tfrac{1}{\eta}\, b^{dis}_{e,t,s} \tag{C-1j}$$

$$SOC_{e,1,s} = 0.5\, \bar{S}_e + \eta\, b^{ch}_{e,1,s} - \tfrac{1}{\eta}\, b^{dis}_{e,1,s} \tag{C-1k}$$

$$SOC_{e,T^{max},s} = 0.5\, \bar{S}_e \tag{C-1l}$$

with bounds $b^{ch}_{e,t,s}, b^{dis}_{e,t,s} \leq \bar{b}_e$ and $SOC_{e,t,s} \leq \bar{S}_e$ (C-1m – C-1o).

A.2. Step 1: Crowley Adding FCR-D up

Crowley introduces FCR-D by replacing the FFR revenue term with an FCR-D revenue term in the objective. The reserved up capacity earns

$$\lambda^{\text{FCRD}}_{t,s}\, y^{acc\text{FCRD}}_{t,s} \sum_{e\in E} p^{\text{FCRD}}_{e,t}, \tag{C-3a}$$

giving the FCR-D objective

$$\begin{aligned} \min \sum_{s\in S} \pi_s \sum_{t\in T} \Big[& \lambda^{im}_{t,s}\, p^{im}_{t,s} - \lambda^{ex}_{t,s} \big(p^{ex}_{t,s} - \sum_{e} p^{\text{FCRD}\uparrow}_{e,t,s}\big) \\ & - \lambda^{\text{FCRD}}_{t,s}\, y^{acc\text{FCRD}}_{t,s} \sum_{e} p^{\text{FCRD}}_{e,t} \Big]. \end{aligned} \tag{C-2a}$$

The acceptance and activation indicators are pre-computed to preserve linearity, and the activation energy is the product

$$p^{\text{FCRD}\uparrow}_{e,t,s} = y^{acc\text{FCRD}}_{t,s}\, y^{act\text{FCRD}}_{t,s}\, p^{\text{FCRD}}_{e,t}. \tag{C-3b}$$

The FCR-D capacity must be deliverable from both a power and a SOC perspective:

$$b^{dis}_{e,t,s} + y^{acc\text{FCRD}}_{t,s}\, p^{\text{FCRD}}_{e,t} \leq \bar{b}_e \tag{C-3c}$$

$$SOC_{e,t,s} \geq y^{acc\text{FCRD}}_{t,s}\, p^{\text{FCRD}}_{e,t} \tag{C-3d}$$

The SOC dynamics are updated to include the FCR-D up activation as additional discharge:

$$SOC_{e,t,s} = SOC_{e,t-1,s} + \eta\, b^{ch}_{e,t,s} - \tfrac{1}{\eta}\big(b^{dis}_{e,t,s} + p^{\text{FCRD}\uparrow}_{e,t,s}\big). \tag{C-2j}$$

A.3. Step 2: Dropping the stochastic case

Crowley's stochastic program is designed to answer the structural question of which products an idealised community would bid into, given a distribution over daily realisations. Our research question is operational: given the best available forecasts for tomorrow, what is the upper-bound revenue that an aggregated BTM portfolio could obtain? This naturally lives in the deterministic regime $|S| = 1$. Sticking with the deterministic case also avoids conflating gains from optimisation with gains from risk-hedging, and lets the deterministic-perfect-forecast outcome serve as a clean upper bound for any practical day-ahead operation. We therefore drop the scenario index s in all variables and parameters.

A.4. Step 3: Calculating the activation share $y_t^{act\mathrm{FCRD}}$

Crowley approximates the within-hour activation share with a linear interpolation $\frac{49.9\,\mathrm{Hz} - f_{t,s}}{0.4\,\mathrm{Hz}}$ based on a single representative frequency reading per hour. Because we obtain Nordic frequency data at 0.1-second resolution we can replace this with the exact empirical share

$$y_t^{act\mathrm{FCRD}\uparrow} = \frac{\#\{\tau \in t : f_\tau < 49.9\,\mathrm{Hz}\}}{\#\{\tau \in t\}}, \tag{12}$$

and analogously for FCR-D down with the $> 50.1\,\mathrm{Hz}$ threshold. As $y_t^{act\mathrm{FCRD}}$ is precomputed and enters only as a coefficient on a decision variable, linearity is preserved.

A.5. Step 4: Calculating the acceptance indicator $y_t^{acc\mathrm{FCRD}}$

Crowley defines $y_{t,s}^{acc\mathrm{FCRD}}$ as a Boolean indicating whether a community's bid price clears the market. The trick that keeps this useful in a program that itself chooses the bid volume is that $y_{t,s}^{acc\mathrm{FCRD}}$ only appears multiplied with the bid volume $p_{e,t}^{\mathrm{FCRD}}$, so an hour with $y_{t,s}^{acc\mathrm{FCRD}} = 0$ is simply excluded from bidding.

In our case study the batteries are pre-installed and financially motivated by self consumption. The marginal participation cost is therefore essentially limited to battery degradation, which for FCR-D's small and infrequent activations is negligible. Because we are computing an upper bound and a rational bidder always bids at or below the realised clearing price, acceptance is by construction. We therefore set $y_t^{acc\mathrm{FCRD}\uparrow} = y_t^{acc\mathrm{FCRD}\downarrow} = 1$ for all hours with non-negative prices, which for 2025 is all hours. The battery constraints (C-3c) and (C-3d) simplify, with $y_t^{acc} = 1$, to (C-3c') and (C-3d'):

$$b_{e,t}^{dis} + p_{e,t}^{\mathrm{FCRD}} \le \bar{b}_e \tag{C-3c'}$$

$$SOC_{e,t} \ge p_{e,t}^{\mathrm{FCRD}}. \tag{C-3d'}$$

A.6. Step 5: Calculating the FCR-D price $\lambda_t^{\mathrm{FCRD}}$

Crowley does not explicitly state how the two FCR-D auction prices are combined into a single $\lambda_t^{\mathrm{FCRD}}$. One could hypothesise that either the highest price is selected of early and late or else that only early participation is taken into account. Since we now use two auction models we use the real prices for late and early instead.

A.7. Step 6: Per-building power balance for BTM aggregation

Where Crowley aggregates assets within a CEC behind a single transformer station (so power can be shared inside the community at no penalty), the BTM setting forbids sharing power between buildings without incurring full tax and tariff. Each building therefore has its own physical grid connection and must satisfy its own power balance. We split the aggregate power balance (C-1b) into per-community balances:

$$p_{e,t} - p_{e,t}^{im} + p_{e,t}^{ex} = 0, \quad \forall e \in E, t \in T \tag{C-1b´}$$

with $p_{e,t}^{im}, p_{e,t}^{ex} \ge 0$. Note that $p_{e,t} = p_{e,t}^{im} - p_{e,t}^{ex}$ is now redundant, but is kept for notational similarity with Crowley. Since we only consider battery prosumers, the sets N and C collapse to E, and the objective is rewritten by summing imports and exports over E:

$$\begin{aligned} \min \sum_{t \in T} \Big[& \lambda_t^{im} \sum_e p_{e,t}^{im} - \lambda_t^{ex} \big(\sum_e p_{e,t}^{ex} - \sum_e p_{e,t}^{\mathrm{FCRD}\uparrow} \big) \\ & - \lambda_t^{\mathrm{FCRD}} \, y_t^{acc\mathrm{FCRD}} \sum_e p_{e,t}^{\mathrm{FCRD}} \Big]. \end{aligned} \tag{C-2a´}$$

A.8. Step 7: Per-battery efficiency in the SOC equations

For a realistic aggregation we let each BTM have its own one-way round-trip efficiency η_e rather than a global η:

$$SOC_{e,t} = SOC_{e,t-1} + \eta_e\, b^{ch}_{e,t} - \tfrac{1}{\eta_e}\big(b^{dis}_{e,t} + p^{\mathrm{FCRD}\uparrow}_{e,t}\big). \tag{C-2j´}$$

A.9. Step 8: LER classification and sustain feasibility

Our aggregated battery portfolio is in essence a battery energy storage system (BESS) as described by Thingvad et al. (2024), with some additional complexities specific to the BTM aggregation setting. It is also definitively a Limited Energy Reservoir (LER) asset. Since we wish to bid approximately our full 1 MW rated capacity into the market, escaping LER classification would require sustaining full output for four consecutive hours — demanding 4 MWh of usable storage. With only 2 MWh at full SOC, we are well within the LER definition (Energinet, 2023).

Drawing upon Thingvad et al. (2024), we integrate their LER requirements into our model. The first requirement concerns how much energy must be held in reserve to sustain output: the regulatory minimum is $\upsilon^D_E = 0.33\,\mathrm{h}$, meaning the battery must be able to sustain full reserved output for $\frac{1}{3}$ h (20 minutes) (Energinet, 2023). Translating Thingvad et al.'s energy constraints (Eqs. 12–13) into our notation and retaining only FCR-D terms gives the simplified form:

$$SOC_{e,t} \geq p^{\mathrm{FCRD}\uparrow res}_{e,t} \cdot \upsilon^D_E, \quad \forall e \in E, t \in T \tag{T-12′}$$

$$\bar{S}_e - SOC_{e,t} \geq p^{\mathrm{FCRD}\downarrow res}_{e,t} \cdot \upsilon^D_E, \quad \forall e \in E, t \in T \tag{T-13′}$$

We extend these with two details absent from Thingvad et al.'s formulation: we condition on bid acceptance via y^{acc}_t, and we account for per-community battery efficiency η_e. For discharge (up direction) the efficiency loss means the battery must hold *more* energy than the raw kWh figure suggests, so we divide by η_e; for charging (down direction) not all absorbed power is stored, so we multiply by η_e. The resulting updated sustain constraints are:

$$SOC_{e,t} \geq y^{acc\mathrm{FCRD}\uparrow}_t \,\frac{\upsilon^D_E}{\eta_e}\, p^{\mathrm{FCRD}\uparrow res}_{e,t}, \quad \forall e \in E, t \in T \tag{E-7a}$$

$$\bar{S}_e - SOC_{e,t} \geq y^{acc\mathrm{FCRD}\downarrow}_t \,\upsilon^D_E\, \eta_e\, p^{\mathrm{FCRD}\downarrow res}_{e,t}, \quad \forall e \in E, t \in T \tag{E-7b}$$

These replace the earlier (E-7a) and (E-7b) pair, which used the more conservative $T_{sustain} = 0.5\,\mathrm{h}$. Setting $\upsilon^D_E = 0.33\,\mathrm{h}$ aligns the model with the current regulatory minimum (Energinet, 2023). The previous (C-3d) bounds ($SOC_{e,t} \geq p^{\mathrm{FCRD}\uparrow res}_{e,t}$ and $\bar{S}_e - SOC_{e,t} \geq p^{\mathrm{FCRD}\downarrow res}_{e,t}$) are now redundant since $\upsilon^D_E < 1$ and are dropped.

A.10. Step 9: LER power constraints – NEM cross-direction reservation

The second LER requirement concerns instantaneous power. A battery inverter has a single rated power limit $\bar{b}_e$ (kW), and every commitment the battery makes at any given moment – FCR-D activation, NEM restoration, and arbitrage – claims a slice of this finite capacity. Because SOC can only move in one direction at a time, the inverter operates either in discharge mode or charge mode each hour; the constraints are therefore enforced separately per direction.

For the *discharge/upward* direction, consider what must fit within $\bar{b}_e$ simultaneously: the full reserved FCR-D up capacity $p^{\mathrm{FCRD}\uparrow res}_{e,t}$ (which discharges if activated), arbitrage discharge $b^{dis}_{e,t}$, and , crucially, a fraction υ^D_P of the FCR-D *down* reservation held available for NEM. This NEM headroom is needed because if FCR-D down previously activated and charged the battery too full, the inverter must be able to discharge at up to $\upsilon^D_P \cdot p^{\mathrm{FCRD}\downarrow res}_{e,t}$ to restore SOC. Symmetrically, for the *charge/downward* direction, the full down reservation, arbitrage charging, and a NEM fraction of the up reservation must all fit within $\bar{b}_e$.

Translating Thingvad et al.'s power constraints (Eqs. 10–11) into our notation, dropping FCR-N and FFR terms, and applying the constraint per community rather than at the aggregate level, yields:

$$\bar{b}_e \geq p^{\mathrm{FCRD}\uparrow res}_{e,t} + p^{\mathrm{FCRD}\downarrow res}_{e,t} \cdot \upsilon^D_P + b^{dis}_{e,t}, \quad \forall e \in E, t \in T \tag{E-8a}$$

$$\bar{b}_e \geq p^{\mathrm{FCRD}\downarrow res}_{e,t} + p^{\mathrm{FCRD}\uparrow res}_{e,t} \cdot \upsilon^D_P + b^{ch}_{e,t}, \quad \forall e \in E, t \in T \tag{E-8b}$$

where $\upsilon_P^D = 0.20$. These replace the simpler constraints from Crowley's formulation, which lacked the cross-direction NEM term. The practical consequence is that when bidding symmetrically in both directions with no arbitrage, the binding constraint becomes $1.2\,p \leq \bar{b}_e$, capping the simultaneous bid at approximately 833 kW per direction rather than the full 1000 kW – consistent with Thingvad et al. (2024), Table I.

The per-community application is correct because each BTM community has its own physical inverter; communities cannot share inverter headroom across buildings. This is a direct consequence of the BTM power balance constraint (C-1b') introduced in Step 6.

A.11. Step 10: Minimum bid size and the move to MILP

A standalone BTM cannot realistically meet the FCR-D minimum bid size on its own, which is precisely the motivation for aggregation. We must therefore constrain the *portfolio-level* reserved capacity to either be zero or to lie in $[p^{min}, \bar{P}]$. Introducing a binary indicator $z_t \in \{0, 1\}$ gives

$$p^{min}\, z_t \leq \sum_{e \in E} p_{e,t}^{\mathrm{FCRD}} \leq \bar{P}\, z_t, \quad z_t \in \{0, 1\}. \tag{E-10}$$

This is the only source of non-linearity in the program and forces the move from LP to MILP. Note that z_t lives only at the portfolio level; we do not constrain individual BTM contributions.

A.12. Step 11: Adding FCR-D down

Where FCR-D up corresponds to outputting power onto the grid to raise frequency, FCR-D down corresponds to absorbing power to lower frequency. For a battery this means the down activation *charges* the battery rather than discharging it, which has implications for energy accounting:

- Activated down energy is absorbed from the grid. To avoid double-counting against import cost (it is already paid for via the reservation revenue, and the BSP role allows only reservation payments), we net activated down energy out of imports in the objective.
- The SOC equation gains a $+\eta_e\, p_{e,t}^{\mathrm{FCRD}\downarrow act}$ term on the charging side.
- The LER power constraint (E-8b) already accounts for the down reservation and its cross-direction NEM requirement.
- Sustain feasibility flips: there must be *free* SOC capacity rather than stored energy, giving $\bar{S}_e - SOC_{e,t} \geq y_t^{acc\mathrm{FCRD}\downarrow}\, T_{sustain}\, \eta_e\, p_{e,t}^{\mathrm{FCRD}\downarrow res}$.
- The minimum bid size constraint is duplicated for the down direction with its own binary $z_t^{\downarrow}$.

The activation linking equation mirrors the up case:

$$p_{e,t}^{\mathrm{FCRD}\downarrow act} = y_t^{acc\mathrm{FCRD}\downarrow}\, y_t^{act\mathrm{FCRD}\downarrow}\, p_{e,t}^{\mathrm{FCRD}\downarrow res}. \tag{E-5b}$$

This completes the derivation of the full early MILP.

A.13. Note on the T^{max} exclusion in (E-6a)

Similar update to (C-2i): $p_{e,t}^{FCRD\uparrow act}$ is added to battery drain and $p_{e,t}^{FCRD\downarrow act}$ to battery gain. Crowley et al. (2024) excludes T^{max} here, applying the constraint only to $t \in T \setminus \{1, T^{max}\}$. In our model this causes a discharge spike at the last hour: the battery sells all remaining SOC because it automatically returns to $0.5\,\bar{S}_e$ at T^{max} without any energy cost. For a revenue calculation this is unrealistic, and to correct this we include T^{max} in the domain, applying the constraint to $t \in T \setminus \{1\}$ instead, forcing the last hour to be subject to the same SOC dynamics as all other hours.

$$SOC_{e,t} = SOC_{e,t-1} + \eta_e\Big(b_{e,t}^{ch} + p_{e,t}^{FCRD\downarrow act}\Big) - \frac{1}{\eta_e}\Big(b_{e,t}^{dis} + p_{e,t}^{FCRD\uparrow act}\Big), \quad \forall e \in E,\ t \in T \setminus \{1\} \tag{E-6a}$$

B. Derivation of the Late FCR-D Model

While many studies simplify FCR-D bidding to participation in one auction ; for instance Thingvad et al. (2024) acknowledge the two-auction structure but assume all capacity is sold on the early auction; in practice the two auctions occur sequentially: the early auction clears at 00:30 the day before delivery, the day-ahead spot price is then published, and the late auction closes at 18:00. The late model treats the second auction explicitly given the cleared early position.

B.1. Assuming the two auctions can be solved separately

A fully optimal joint formulation would solve both auctions simultaneously while respecting that the late auction's information set strictly contains the early's. Such a formulation would, for example, allow strategic over- or under-bidding in the early auction in anticipation of the cheapest possible late adjustment, and would require modelling the distribution of late prices and forecasts at early-bid time. This is genuinely two-stage stochastic and quickly becomes difficult to handle for the daily horizon we are interested in. Moreover, the relative ordering of early and late prices is not systematic in the data: it is not always profitable to over- or under-bid in either direction. We therefore make the simplifying assumption that the early auction is solved independently with the best available early-time forecasts, and that its accepted volume $p_{e,t}^{\mathrm{FCRD}\uparrow acc}, p_{e,t}^{\mathrm{FCRD}\downarrow acc}$ is taken as fixed input to the late problem. This does not guarantee mathematical optimality across the joint problem but is a reasonable starting point that isolates the contribution of late-auction information, and the looseness can be improved later by improving early-time forecast accuracy.

B.2. Late-stage decision variables: top-ups and buy-backs

At the late auction the aggregator has two ways to change its position relative to the early-cleared volume:

1. Place additional positive bids ($a_{e,t}^{\uparrow}, a_{e,t}^{\downarrow}$): top up the reserved capacity for that hour and direction.
2. Place negatively-oriented bids ($c_{e,t}^{\uparrow}, c_{e,t}^{\downarrow}$): cancel part or all of the early position by "buying back" capacity from the market.

The final reserved capacity, the quantity used by all physical/battery constraints, is then

$$p_{e,t}^{\mathrm{FCRD}\uparrow res} = p_{e,t}^{\mathrm{FCRD}\uparrow acc} + a_{e,t}^{\uparrow} - c_{e,t}^{\uparrow},$$

and analogously for the down direction.

B.3. Pricing of top-ups and buy-backs

A top-up earns the late price λ_t^{late}, adding $-\lambda_t^{late}\, a_{e,t}$ to the cost objective. A buy-back however is done at the maximum price of late and early, ensuring we cannot leverage this difference for arbitrage:

$$\lambda_t^{\mathrm{FCRD}\uparrow bb} = \max\{\lambda_t^{\mathrm{FCRD}\uparrow early},\ \lambda_t^{\mathrm{FCRD}\uparrow late}\}.$$

So in our objective we have the accepted bid vector from early which we modify for each hour by adding the cost of cancelling bids and the additional revenue for topping up.

B.4. Minimum bid size constraints on top-ups and finals

Two minimum-bid constraints are needed:

1. Any new positive volume submitted to the late auction must itself satisfy the minimum bid floor; this is the role of w_t and gives constraint (L-2a)–(L-2c).
2. The final reserved quantity (after both auctions) must *also* be either zero or above p^{min}. Otherwise a partial cancellation could leave a stranded position in the open range $(0, p^{min})$, which is infeasible from the TSO's perspective. This is the role of z_t in (L-3a)–(L-3c).

Having both constraints means we cannot in any direction end up with $0 <$ final $< p^{min}$: either cancel all the way to zero, or end with at least the minimum bid.

B.5. No simultaneous top-up and buy-back

Topping up and buying back the same hour in the same direction is meaningless: it would amount to round-tripping volume through the market, paying spreads/fees, without changing the final position. We enforce this with the constraint

$$c_{e,t}^{\uparrow} \le p_{e,t}^{\mathrm{FCRD}\uparrow acc}\,(1 - w_t^{\uparrow})$$

and analogously for the down direction. When $w_t^{\uparrow} = 1$ (we are topping up that hour and direction), the right-hand side collapses to zero and no buy-back is permitted; when $w_t^{\uparrow} = 0$, the buy-back is upper-bounded by what was accepted in the early auction (you cannot cancel more than you hold). This is a linear constraint because it relies on the existing binary $w_t^{\uparrow}$.

B.6. Carried-over constraints

All physical constraints from the early model continue to hold using the linked $p_{e,t}^{\mathrm{FCRD}\uparrow res}$ and $p_{e,t}^{\mathrm{FCRD}\downarrow res}$ defined by (L-1a) and (L-1b). In particular the activation tracking (E-5a), SOC dynamics (E-6a)–(E-6c), sustain feasibility (E-7a) and (E-7b), LER power limits (E-8a), (E-8b), and battery capacity limits are all kept verbatim – they constrain the final reserved capacity and so are agnostic to whether that capacity came from the early auction, top-ups, or net of buy-backs. This concludes the derivation of the late MILP.

C. Complete Equation Listing: Early Auction MILP

For reference, we list every equation in the early auction model in one place. Sets: E (energy communities), T (hourly time steps). Parameters: $\upsilon_P^D = 0.20$ (NEM power reservation factor), $\upsilon_E^D = 0.33\,\mathrm{h}$ (sustain duration, regulatory minimum of 20 minutes).

Objective function:

$$\min \sum_{t\in T} \left[\lambda_t^{im} \left(\sum_e p_{e,t}^{im} - \sum_e p_{e,t}^{\mathrm{FCRD}\downarrow act} \right) - \lambda_t^{ex} \left(\sum_e p_{e,t}^{ex} - \sum_e p_{e,t}^{\mathrm{FCRD}\uparrow act} \right) - \lambda_t^{\mathrm{FCRD}\uparrow} \sum_e p_{e,t}^{\mathrm{FCRD}\uparrow res} - \lambda_t^{\mathrm{FCRD}\downarrow} \sum_e p_{e,t}^{\mathrm{FCRD}\downarrow res} \right] \tag{E-obj}$$

Per-community power balance:

$$p_{e,t} - p_{e,t}^{im} + p_{e,t}^{ex} = 0, \qquad \forall e \in E,\, t \in T \tag{E-1}$$
$$p_{e,t}^{im} \geq 0, \qquad \forall e \in E,\, t \in T \tag{E-2}$$
$$p_{e,t}^{ex} \geq 0, \qquad \forall e \in E,\, t \in T \tag{E-3}$$

Battery power balance (net consumption):

$$p_{e,t} + PV_{e,t} - D_{e,t} - \left(b_{e,t}^{ch} + p_{e,t}^{\mathrm{FCRD}\downarrow act}\right) + \left(b_{e,t}^{dis} + p_{e,t}^{\mathrm{FCRD}\uparrow act}\right) = 0, \quad \forall e \in E,\, t \in T \tag{E-4}$$

FCR-D activation tracking:

$$p_{e,t}^{\mathrm{FCRD}\uparrow act} = y_t^{act\mathrm{FCRD}\uparrow}\, p_{e,t}^{\mathrm{FCRD}\uparrow res}, \qquad \forall e \in E,\, t \in T \tag{E-5a}$$
$$p_{e,t}^{\mathrm{FCRD}\downarrow act} = y_t^{act\mathrm{FCRD}\downarrow}\, p_{e,t}^{\mathrm{FCRD}\downarrow res}, \qquad \forall e \in E,\, t \in T \tag{E-5b}$$

State of charge dynamics:

$$SOC_{e,t} = SOC_{e,t-1} + \eta_e\left(b_{e,t}^{ch} + p_{e,t}^{\mathrm{FCRD}\downarrow act}\right) - \tfrac{1}{\eta_e}\left(b_{e,t}^{dis} + p_{e,t}^{\mathrm{FCRD}\uparrow act}\right), \qquad \forall e \in E,\, t \in T \setminus \{1\} \tag{E-6a}$$
$$SOC_{e,1} = 0.5\,\bar{S}_e + \eta_e\left(b_{e,1}^{ch} + p_{e,1}^{\mathrm{FCRD}\downarrow act}\right) - \tfrac{1}{\eta_e}\left(b_{e,1}^{dis} + p_{e,1}^{\mathrm{FCRD}\uparrow act}\right), \qquad \forall e \in E \tag{E-6b}$$
$$SOC_{e,T^{max}} = 0.5\,\bar{S}_e, \qquad \forall e \in E \tag{E-6c}$$

Sustain feasibility (LER energy constraints):

$$SOC_{e,t} \geq \frac{\upsilon_E^D}{\eta_e}\, p_{e,t}^{\mathrm{FCRD}\uparrow res}, \qquad \forall e \in E,\, t \in T \tag{E-7a}$$
$$\bar{S}_e - SOC_{e,t} \geq \upsilon_E^D\, \eta_e\, p_{e,t}^{\mathrm{FCRD}\downarrow res}, \qquad \forall e \in E,\, t \in T \tag{E-7b}$$

LER power constraints (NEM cross-direction reservation):

$$b_{e,t}^{dis} + p_{e,t}^{\mathrm{FCRD}\uparrow res} + \upsilon_P \cdot p_{e,t}^{\mathrm{FCRD}\downarrow res} \leq \bar{b}_e, \qquad \forall e \in E,\, t \in T \tag{E-8a}$$
$$b_{e,t}^{ch} + p_{e,t}^{\mathrm{FCRD}\downarrow res} + \upsilon_P \cdot p_{e,t}^{\mathrm{FCRD}\uparrow res} \leq \bar{b}_e, \qquad \forall e \in E,\, t \in T \tag{E-8b}$$

Minimum bid size (portfolio level):

$$p^{min}\, z_t^{\uparrow} \leq \sum_{e \in E} p_{e,t}^{\mathrm{FCRD}\uparrow res} \leq \bar{P}\, z_t^{\uparrow}, \qquad \forall t \in T \tag{E-10a}$$

$$p^{min}\, z_t^{\downarrow} \leq \sum_{e \in E} p_{e,t}^{\mathrm{FCRD}\downarrow res} \leq \bar{P}\, z_t^{\downarrow}, \qquad \forall t \in T \tag{E-10b}$$

$$z_t^{\uparrow},\, z_t^{\downarrow} \in \{0, 1\}, \qquad \forall t \in T \tag{E-10c}$$

Non-negativity:

$$p_{e,t}^{\mathrm{FCRD}\uparrow res},\ p_{e,t}^{\mathrm{FCRD}\downarrow res},\ b_{e,t}^{ch},\ b_{e,t}^{dis},\ SOC_{e,t} \geq 0, \quad \forall e \in E,\, t \in T \tag{E-11}$$

D. Complete Equation Listing: Late Auction MILP

The late auction model takes the accepted early bids $p_{e,t}^{\mathrm{FCRD}\uparrow acc}$ and $p_{e,t}^{\mathrm{FCRD}\downarrow acc}$ as given parameters and introduces top-up and buy-back decision variables.

Linking equations (final reserved = early accepted + top-up − buy-back):

$$p_{e,t}^{\mathrm{FCRD}\uparrow res} = p_{e,t}^{\mathrm{FCRD}\uparrow acc} + a_{e,t}^{\uparrow} - c_{e,t}^{\uparrow}, \qquad \forall e \in E,\, t \in T \tag{L-1a}$$

$$p_{e,t}^{\mathrm{FCRD}\downarrow res} = p_{e,t}^{\mathrm{FCRD}\downarrow acc} + a_{e,t}^{\downarrow} - c_{e,t}^{\downarrow}, \qquad \forall e \in E,\, t \in T \tag{L-1b}$$

Objective function:

$$\begin{aligned}\min \sum_{t\in T} \Bigg[& \lambda_t^{im}\left(\sum_e p_{e,t}^{im} - \sum_e p_{e,t}^{\mathrm{FCRD}\downarrow act}\right) - \lambda_t^{ex}\left(\sum_e p_{e,t}^{ex} - \sum_e p_{e,t}^{\mathrm{FCRD}\uparrow act}\right) \\ & - \lambda_t^{\mathrm{FCRD}\uparrow early}\sum_e p_{e,t}^{\mathrm{FCRD}\uparrow acc} - \lambda_t^{\mathrm{FCRD}\downarrow early}\sum_e p_{e,t}^{\mathrm{FCRD}\downarrow acc} \\ & + \lambda_t^{\mathrm{FCRD}\uparrow bb}\sum_e c_{e,t}^{\uparrow} + \lambda_t^{\mathrm{FCRD}\downarrow bb}\sum_e c_{e,t}^{\downarrow} \\ & - \lambda_t^{\mathrm{FCRD}\uparrow late}\sum_e a_{e,t}^{\uparrow} - \lambda_t^{\mathrm{FCRD}\downarrow late}\sum_e a_{e,t}^{\downarrow}\Bigg]\end{aligned} \tag{L-obj}$$

where $\lambda_t^{\mathrm{FCRD}\uparrow bb} = \max\{\lambda_t^{\mathrm{FCRD}\uparrow early},\, \lambda_t^{\mathrm{FCRD}\uparrow late}\}$ and analogously for down.

Minimum bid size for top-ups:

$$p^{min}\, w_t^{\uparrow} \le \sum_e a_{e,t}^{\uparrow} \le \bar{P}\, w_t^{\uparrow}, \qquad \forall t \in T \tag{L-2a}$$

$$p^{min}\, w_t^{\downarrow} \le \sum_e a_{e,t}^{\downarrow} \le \bar{P}\, w_t^{\downarrow}, \qquad \forall t \in T \tag{L-2b}$$

$$w_t^{\uparrow},\, w_t^{\downarrow} \in \{0,1\}, \qquad \forall t \in T \tag{L-2c}$$

Minimum bid size for final reservation:

$$p^{min}\, z_t^{\uparrow} \le \sum_e p_{e,t}^{\mathrm{FCRD}\uparrow res} \le \bar{P}\, z_t^{\uparrow}, \qquad \forall t \in T \tag{L-3a}$$

$$p^{min}\, z_t^{\downarrow} \le \sum_e p_{e,t}^{\mathrm{FCRD}\downarrow res} \le \bar{P}\, z_t^{\downarrow}, \qquad \forall t \in T \tag{L-3b}$$

$$z_t^{\uparrow},\, z_t^{\downarrow} \in \{0,1\}, \qquad \forall t \in T \tag{L-3c}$$

No simultaneous top-up and buy-back:

$$c_{e,t}^{\uparrow} \le p_{e,t}^{\mathrm{FCRD}\uparrow acc}\,(1 - w_t^{\uparrow}), \qquad \forall e \in E,\, t \in T \tag{L-4a}$$

$$c_{e,t}^{\downarrow} \le p_{e,t}^{\mathrm{FCRD}\downarrow acc}\,(1 - w_t^{\downarrow}), \qquad \forall e \in E,\, t \in T \tag{L-4b}$$

Carried-over constraints from the early model:

All physical constraints from the early model carry over with $p_{e,t}^{\mathrm{FCRD}\uparrow res}$ and $p_{e,t}^{\mathrm{FCRD}\downarrow res}$ now defined by the linking equations (L-1a) and (L-1b):

- Per-community power balance: (E-1)–(E-3)
- Battery power balance: (E-4)
- Activation tracking: (E-5a), (E-5b)
- State of charge dynamics: (E-6a)–(E-6c)
- Sustain feasibility: (E-7a), (E-7b)
- LER power constraints: (E-8a), (E-8b)
- Non-negativity of $p_{e,t}^{\mathrm{FCRD}\uparrow res}$, $p_{e,t}^{\mathrm{FCRD}\downarrow res}$, $b_{e,t}^{ch}$, $b_{e,t}^{dis}$, $SOC_{e,t}$, $a_{e,t}^{\uparrow}$, $a_{e,t}^{\downarrow}$, $c_{e,t}^{\uparrow}$, $c_{e,t}^{\downarrow} \geq 0$

Deriving Equations Overview

The table below traces every equation in the final early and late models back to its origin. Prefixes: **C** = Crowley et al. (2024) original; **T** = Thingvad et al. (2024) original; **E** = final early auction model; **L** = final late auction model. A prime ($'$) denotes one modification from the source equation.

Final label	Derivation chain	What changed at each step
	Objective	
E-obj	C-1a → C-2a → C-2a$'$ → E-obj	Base import/export cost; add FCR-D up revenue; drop stochastic index, sum over communities; add FCR-D down revenue and netting
Power balance		
E-1	C-1b → C-1b$'$ → E-1	Aggregate community balance; split to per-building BTM balance
E-2, E-3	C-1c, C-1d → E-2, E-3	Unchanged non-negativity of import and export
E-4	C-1i → E-4	Battery power balance extended with FCR-D up and down activation flows
Activation tracking		
E-5a	C-3b → E-5a	FCR-D up activation linking; y^{acc} set to 1 since acceptance by construction
E-5b	C-3b (new for down) → E-5b	New equation: mirrors C-3b for FCR-D down direction
State of charge dynamics		
E-6a	C-1j → C-2j → C-2j$'$ → E-6a	Global η; add FCR-D up discharge; per-community η_e; add FCR-D down charging
E-6b	C-1k → E-6b	Initial SOC boundary extended with FCR-D activation flows and per-community η_e
E-6c	C-1l → E-6c	Terminal SOC boundary, unchanged
Sustain feasibility (LER energy)		
E-7a	C-3d → C-3d$'$ → T-12$'$ → E-7a	SOC $\geq y^{acc}p^{res}$; set $y^{acc} = 1$; replace 1 h with $\upsilon_E^D = 0.33$ h (Thingvad); add per-community η_e
E-7b	T-13$'$ → E-7b	New from Thingvad for down direction; add η_e and y^{acc}
LER power constraints		
E-8a	T-10 → E-8a	Thingvad discharge budget; drop FCR-N and FFR terms; apply per-community
E-8b	C-3c → C-3c$'$ → T-11 → E-8b	Crowley charge limit; set $y^{acc} = 1$; add 20% NEM cross-direction term (Thingvad); apply per-community
SOC capacity limits		
E-9a, E-9b	C-3d → C-3d$'$ → *dropped*	Crowley SOC lower bound (implicit 1 h sustain); set $y^{acc} = 1$; dropped in favour of (E-7a)–(E-7b) which encode the regulatory minimum of 20 minutes — retaining both would leave (E-7a)–(E-7b) permanently slack
Minimum bid and non-negativity		
E-10a,b,c	New → E-10a,b,c	No Crowley equivalent; portfolio-level minimum bid with binary z_t
E-11	C-1m, C-1n → E-11	Extended to cover FCR-D reservation and activation variables

Data-set Overview

A closer overview of data-set columns and how they link to model variables as well as a brief description to complement existing tables under the data handling section.

Table 10. Overview of the dataset columns used in the model, grouped by purpose and mapped to the corresponding model parameter. All prices are in Øre/kWh, note later converted to DKK/kWh, energy in kWh/h, and frequency in Hz. The full dataset contains 17,520 rows (8,760 hours × 2 EC profile types).

Column name	**Model symbol**	**Unit**	**Source / description**
Index & identification			
`hour_utc`	$t \in T$	datetime	Hourly UTC timestamp, full year 2025
`ec_id`	$e \in E$	$\{\mathrm{b}, \mathrm{s}\}$	EC profile type; two real communities (b, s) used to construct 10 simulated ECs
`data_quality_flag`	—	$\{0, 1\}$	Flag for linearly interpolated null values
Energy community profiles			
`consumption`	$D_{e,t}$	kWh/h	Aggregate hourly load per community; Enyday A/S private dataset
`pv_production_kwh`	$PV_{e,t}$	kWh/h	Aggregate hourly PV generation per community; Enyday A/S private dataset
Spot price & constructed energy prices			
`spot_exkl_vat_ore_kwh`	P_t^{spot}	Øre/kWh	Day-ahead spot price DK2; NordPool via Energi Data Service (hourly Jan–Sep, 15-min avg Oct–Dec)
`buy_price_inkl_vat_ore_kwh`	λ_t^{im}	Øre/kWh	Import price incl. DSO tariff (Radius), state fees, VAT; Eq. (1)
`sell_price_inkl_vat_ore_kwh`	λ_t^{ex}	Øre/kWh	Export price minus feed-in tariff; Eq. (2)
FCR-D reservation prices			
`price_..fcr_d_ned_..._early`	$\lambda_t^{\text{FCR-D}\downarrow\text{early}}$	Øre/kWh	FCR-D down early auction clearing price; Energi Data Service
`price_..fcr_d_ned_..._late`	$\lambda_t^{\text{FCR-D}\downarrow\text{late}}$	Øre/kWh	FCR-D down late auction clearing price; Energi Data Service
`price_..fcr_d_upp_..._early`	$\lambda_t^{\text{FCR-D}\uparrow\text{early}}$	Øre/kWh	FCR-D up early auction clearing price; Energi Data Service
`price_..fcr_d_upp_..._late`	$\lambda_t^{\text{FCR-D}\uparrow\text{late}}$	Øre/kWh	FCR-D up late auction clearing price; Energi Data Service
Grid frequency & activation fractions			
`freq_avg_hz`	—	Hz	Hourly average Nordic grid frequency; Fingrid dataset 339 (0.1 s resolution, aggregated)
`y_act_fcrd_up`	$y_t^{\text{actFCRD}\uparrow}$	frac.	Fraction of hour with $f < 49.9\,\text{Hz}$; pre-computed FCR-D up activation share
`y_act_fcrd_down`	$y_t^{\text{actFCRD}\downarrow}$	frac.	Fraction of hour with $f > 50.1\,\text{Hz}$; pre-computed FCR-D down activation share