energies

MDPI

Article

A Real-Life Demonstration of Secondary Frequency Reserve Provision with Electric Water Heaters

Louis Brouyaux [1,*], Sandro Iacovella [2] and Sylvain Quoilin [1,*]

[1] Integrated & Sustainable Energy Systems, University of Liège, 7 Place du Vingt Août, 4000 Liège, Belgium
[2] ThermoVault BV, 32/3A Diestsevest, 3000 Leuven, Belgium
* Correspondence: louis.brouyaux@student.uliege.be (L.B.); squoilin@uliege.be (S.Q.)

Abstract: Residential electric water heaters have the potential to significantly contribute to the balancing of the grid by providing frequency services. However, this entails a large-scale, challenging control problem subject to several uncertainties. In this paper, we perform the first real-life validation of secondary frequency reserve provision with a cluster of residential thermal loads in a near-commercial setting. We adopt an aggregate-and-dispatch control approach, which combines a scalable optimization step enabled by a reduced-order model with a real-time dispatch step. To handle the uncertainty related to service activation, we incorporate chance constraints in the optimization model and reformulate it as a robust problem. We validate the control approach under the assumption of perfect merit order knowledge in different stages, with a cluster of up to 600 electric water heaters, and show that this pool is able to effectively provide reserves, and that the integration of the chance constraints is beneficial for performance.

Keywords: demand response; ancillary services; frequency control; water heating; virtual power plant; model predictive control; chance constraints; optimization; experimental validation

Academic Editor: Seung-Hoon Yoo

Received: 14 February 2025
Revised: 12 March 2025
Accepted: 24 March 2025
Published: 28 March 2025

Citation: Brouyaux, L.; Iacovella, S.; Quoilin, S. A Real-Life Demonstration of Secondary Frequency Reserve Provision with Electric Water Heaters. *Energies* **2025**, *18*, 1704. https://doi.org/10.3390/en18071704

1. Introduction

1.1. Motivation

The transition to a low-carbon energy system relies on integrating renewable energy sources, such as wind and solar power. However, these sources are inherently variable and unpredictable, creating challenges in maintaining a stable and reliable power grid. To address this, flexible technologies that can adjust their electricity consumption in response to grid needs are crucial.

Electric storage Water Heaters (EWHs) are one such flexible technology. They are widely used in residential homes, with their hot water tanks acting as energy buffers. This means their energy consumption can be shifted over time, without significantly impacting user comfort [1]. As a result, EWHs present a large, untapped potential to support the power grid through demand response. In 2014, more than 20% of European households relied on EWHs as their primary source of domestic hot water [2], making them a significant asset for grid flexibility.

An aggregator could leverage the flexibility of residential EWHs by coordinating a large number of appliances to provide frequency control. Automatic Frequency Restoration Reserves (aFRRs), also called secondary frequency reserves or regulation services, are a good option to valorize EWHs. However, providing aFRRs with EWHs entails several challenges. Firstly, it constitutes a large-scale coordination problem, which can become heavy in terms of computation and communication as the number of steered appliances

grows. Secondly, secondary reserves are procured in competitive markets, and the activated energy volumes depend on variables such as the state of the transmission grid and the bids of other players. The aggregator needs to guarantee a high reliability of service provision, while respecting the technical constraints of each EWH in this uncertain environment. A bidding strategy that is able to anticipate different scenarios of reserve provision is therefore essential.

The objective of this paper was to develop an algorithm that allows an aggregator to provide aFRRs with a cluster of electric water heaters in a realistic market setting with other participants. We also aimed to validate the algorithm with a real-life pool of appliances in an emulated aFRR market, to illustrate its viability and performance in practice.

1.2. Related Work and Contribution

The idea of controlling aggregations of thermal loads to provide ancillary services has been extensively studied in literature. Early works on the topic included [3,4]. More recently, a lot of research has focused on the development of aggregate-and-dispatch strategies to tackle the challenges related to the scale of the aggregate control problem. In these approaches, a reduced-order representation of the cluster of appliances is used to schedule the aggregate power consumption in a computationally efficient way, typically in an optimization model. The aggregate setpoints are then tracked by switching heaters ON and OFF. Aggregate-and-dispatch strategies were, for example, developed in [5,6] for thermal loads and in [7] for electric vehicles. Our work implements a similar approach to achieve scalable control of the pool of EWHs.

Besides theoretical and simulation-based research, there have only been a few works that demonstrated demand response use cases with aggregations of thermal loads in real life. A frequency control pilot with aggregated resources was carried out in 2014 [8]. This work showed how a cluster of heat pumps was able to follow a predefined aggregate power profile. Reference [9] described load reduction experiments carried out with a group of heating and cooling devices in buildings. The baseline consumption, reduction potential, and rebound were predicted based on historical consumption data. A comparable use case was also investigated in real life with residential heat pumps in [10]. The authors made similar predictions using detailed gray-box appliance models. Though both works predicted the response of the pool to a single, specific load reduction request, they did not actively schedule the power consumption of the cluster beyond this event. In contrast, an algorithm for frequency service provision with a broader scope was presented in [11] and tested in real life in [12]. The authors focused on the heating, ventilation, and air conditioning system of a commercial building and designed a comprehensive hierarchical Model Predictive Control (MPC) framework that covered the modeling, scheduling, and real-time control challenges of the problem. Compared to [9,10], such a method allows more precise control of the appliances' power consumption, while accounting for their constraints. It also allows for control and minimization of rebound effects. Therefore, we intend to implement a similar approach for a cluster of EWHs in this work.

To enable aFRR provision with energy-constrained EWHs, we want to develop an optimization-based bidding strategy to hedge against the uncertainty of the aFRR activation volumes, tied to a control signal dispatched by the Transmission System Operator (TSO). One way to achieve this is to rely on scenarios, either drawn from historical data or sampled from representative distributions. In [13], a bidding strategy for secondary reserves based on scenarios of the control signal was proposed. The authors of [14] considered an aggregator of prosumers and used stochastic optimization to handle uncertainties such as household consumption. Alternatively, uncertainties can also be represented as a set,

leading to a robust optimization framework. In [15,16], robust sets were used to ensure the feasibility of a frequency control strategy with thermal loads and photovoltaic inverters.

Bidding strategies that attempt to ensure feasibility for all possible realizations of uncertainty typically tend to produce conservative solutions. A popular alternative is to constrain the feasibility of the schedule with a high probability, resulting in a chance-constrained problem. When they are used to reformulate such probabilistic constraints, scenario-based approaches are notorious for resulting in heavy computations [17]. Conversely, robust approaches can be used together with properly designed uncertainty sets to guarantee constraint satisfaction with a specific confidence level, without increasing the computational complexity of the problem. Such approaches have been shown to perform well in other reserve offering problems [18]. Several options exist to design the uncertainty set, depending on what is known about the uncertain parameters. In [19,20], robust approaches were designed based on full knowledge of the distribution of the uncertainties. When historical data are available but the distribution cannot reliably be inferred, it is possible to represent the uncertainty in a distributionally robust way. This can be based, for example, on empirical statistical parameters [21–24] or Wasserstein ambiguity sets [25]. The set used in [24] is a good option to model the uncertainty of the aFRR control signal, as it was shown to provide one of the least conservative approximations of the desired confidence level among alternatives [26], and it can be constructed efficiently from historical data.

In European markets, aFRR volumes are typically activated following Merit Order (MO), as recommended by European regulation [27]. This ensures that the cheapest energy bids are utilized first, hence minimizing the cost of balancing. However, from the perspective of the aggregator, this means that the activation volumes are heavily dependent on its position in the MO. To capture this influence, the authors of [28,29] proposed discretizing the MO position domain into levels, each associated with a different probability of activation. It is then possible to construct chance constraints based on level-specific random variables [29].

The aggregator's MO position is determined by its bidding price, as well as that of the other market participants. When hedging against activation uncertainty, it is therefore necessary to integrate market information in the scheduling of the aggregator. The authors of [28,29] assumed that the bids of other markets participants were readily available and incorporated that information in their strategy by explicitly modeling the market clearing. The work represented a battery manager as a price maker by formulating a bi-level optimization problem, which embedded the battery's profit maximization, as well as the clearing. An alternative to this approach was proposed in [30]: their work modeled the market with a surrogate neural network, to avoid the computational challenges associated with a bi-level formulation. As the behaviors of other market players are not known in practice, other works have developed methods to predict the MO curve. A line of research achieved this by constructing the Residual Demand Curve (RDC) of the bidder and predicting the RDC based on historical data. The authors of [31] developed a functional prediction framework that produced a deterministic RDC forecast. In [32], historical data were leveraged to generate RDC scenarios, which were used in a stochastic optimization model that maximized profits on a reserves market. In this work, we make the assumption that the aggregator has enough knowledge about the MO curve to link its bidding price to a position.

The main contribution of this work is the first real-life demonstration of secondary frequency reserve provision with an aggregation of thermal loads in a realistic setting. To address the challenges associated with this control problem, we propose an aggregate-and-

dispatch framework adapted from our previous work on primary reserves [33], in which we extend the scheduling step to fit current and future European aFRR markets.

First, we formulate an optimization problem which maximizes the profit of the EWH aggregator, with probabilistic constraints that account for the aFRR activation uncertainty. The aggregator is modeled as an emerging, price-taking player that represents a small share of the aFRR market. To reformulate these constraints in a tractable way, we provide a distributionally robust, probabilistic representation of the aFRR activation uncertainty, based on the forward and backward deviation metrics used in [24,33]. Following the discretization idea of [28,29], we compute uncertainty sets for different MO positions in the aFRR energy auction, based on historical data from the Belgian TSO [34]. These sets allow us to cast our problem as a linear robust optimization model.

The aggregate-and-dispatch control approach with novel optimization model was validated in real life under the assumption that the MO is known. Its performance was assessed based on the bids placed by the optimization, the energy consumption of the individual EWHs, and the reliability of service provision.

The rest of this paper is organized as follows: Section 2 summarizes the aFRR bidding and control problem. Section 3 lays out the aggregate-and-dispatch control structure. Section 4 details the novel scheduling framework. Section 5 presents the experimental setup and the results from the real-life validation. Finally, Section 6 concludes the paper.

2. Problem Statement

Figure 1 positions aFRRs within the balancing services typically used in Europe to maintain grid frequency. Upon a frequency deviation, Frequency Containment Reserves (FCRs) are activated first and have the primary goal of preventing the frequency from deviating further away from its nominal value. aFRRs are activated in a second stage and are used to bring the frequency back to its nominal value. If the aFRRs are not sufficient, they can be complemented by manual Frequency Restoration Reserves (mFRRs) and Replacement Reserves (RRs).

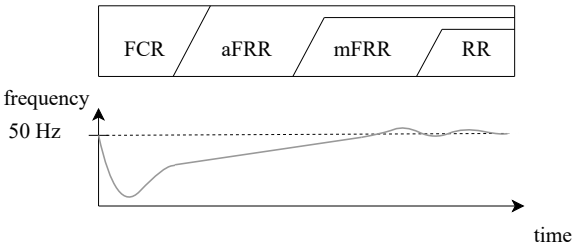

Figure 1. Overview of the balancing services typically used in Europe.

In Europe, secondary reserves are typically procured via two types of auctions. Capacity auctions take place one or several days ahead for time blocks of several hours and are used to ensure the availability of a predefined volume of reserves. These auctions are associated with a remuneration on the awarded capacity. Energy auctions take place closer to real time (the bid submission deadline, or Gate Closing Time (GCT) is 25 min before service delivery in Belgium) for time blocks of 15 min. These auctions are associated with a remuneration on the activated energy. A certain amount of awarded capacity comes with the obligation to bid at least that amount in all the energy auctions taking place during the capacity block. However, it is also possible to bid freely in the energy auctions without participating in the capacity auction. This is less lucrative but results in a lot more flexibility to manage the assets providing the reserve. Therefore, this paper focuses on participation in aFRRs with free energy bids.

During the Delivery Period (DP), the TSO determines the aFRR volume to be activated every 4 s. This volume is activated through the participants of the energy auction, according to the MO. The MO position of the aggregator's bid is determined by its bid price and the bids of the other participants. This is illustrated in Figure 2.

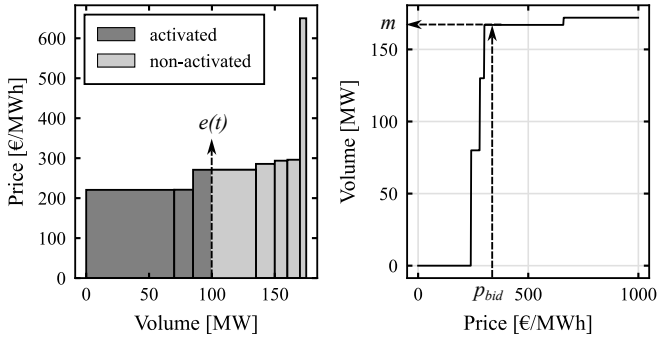

Figure 2. Example of MO for aFRR energy auction in Belgium. **Left**: MO curve and illustration of the MO-based activation of aFRR energy. **Right**: inverted MO curve, which determines the MO position of the aggregator's bid based on its bid price. e is the instantaneous aFRR system activation, m is the merit order position, and p_{bid} is the aFRR bid price.

Reserve providers perform activations by deviating from their baseline consumption, which needs to be communicated to the TSO shortly before delivery (1 min in Belgium). The pilot only covers energy bids and activations associated with a decrease in the consumption compared to the baseline. This is commonly referred to as upward aFRR (aFRR$_{up}$). However, we will refer to this service as aFRR in the rest of the paper for conciseness. It should also be noted that the presented framework can easily be generalized to provide service in both directions.

For the aggregator, the provision of aFRRs with free bids entails three tasks: (i) estimate its bidding capacity and place bids with a specified volume and price 25 min before delivery; (ii) estimate its baseline consumption and declare this consumption 1 min before delivery; (iii) track the power reference resulting from their own baseline and the activation setpoint sent by the TSO every 4 s. This is summarized in Figure 3.

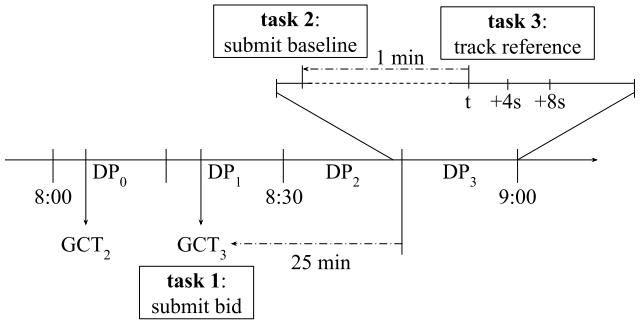

Figure 3. aFRR timings in Belgium (free bidding).

3. Control Architecture

Each EWH used in the pilot was equipped with a control module able to switch the power ON or OFF with a relay. This module also measures the power consumption and water temperature of the appliance, and is able to connect to a cloud computing platform over the internet. The cloud-based central entity is able to gather EWH data in real time and can switch appliances ON and OFF individually. To make sure that the EWH temperature stays within its safety and comfort limits, the following local overrule of the relay state is active:

$$s_i(t) = \begin{cases} 1 \text{ if } T_i(t) \leq T_{min,i}(t) \\ 0 \text{ if } T_i(t) \geq T_{max,i}(t) \end{cases} \tag{1}$$

In the above equation, s is the EWH's ON/OFF state, t is the time, i denotes the EWH index within the pool, T denotes the water temperature, and $[\cdot]_{min,max}$ denotes the minimum and maximum limits of T. Based on this infrastructure, to ensure the scalability of our control with a large number of appliances, we implemented an aggregate-and-dispatch approach. An overview of the control architecture is shown in Figure 4. The methodology consists of three main steps. First, the aggregator uses individual EWH data to create a first-order Virtual Battery Model (VBM) that represents the cluster of devices. To obtain this VBM, we first identify individual EWH models based on the collected data:

$$\begin{aligned} \frac{dT_i(t)}{dt} &= -\alpha_i(T_i(t) - T_{a,i}(t)) \\ &\quad + \beta_i(s_i(t)P_{rated,i} - P_{w,i}(t)) \end{aligned} \tag{2}$$

In the above equation, α and β are the cooling and heating parameters of the thermal model, T_a is the ambient temperature, P_{rated} is the EWH's rated power, and P_w is the equivalent power associated with domestic hot water consumption. The VBM can then be expressed as

$$\frac{dx(t)}{dt} = -\alpha x(t) + P(t) - P_w(t) \tag{3a}$$

$$\alpha = \frac{1}{n_{ewh}} \sum_{i=1}^{n_{ewh}} \alpha_i \tag{3b}$$

$$x(t=0) = \sum_{i=1}^{n_{ewh}} \frac{T_i(t=0) - T_{a,i}(t=0)}{\beta_i} \tag{3c}$$

$$P_w(t) = \sum_{i=1}^{n_{ewh}} P_{w,i}(t) \tag{3d}$$

$$P_{min}(t) \leq P(t) \leq P_{max}(t) \tag{3e}$$

$$x_{min}(t) \leq x(t) \leq x_{max}(t) \tag{3f}$$

In this model, x denotes the VBM state of charge, P denotes the pool's aggregate power consumption, and n_{ewh} is the number of EWHs in the pool. The power and energy limits of Equations (3e) and (3f) need to be chosen so that system trajectories deemed feasible by the VBM are also feasible for the pool of appliances, without violating their individual constraints. Several techniques have been proposed in the literature to determine these bounds. Reference [35] provided a comparison of the main existing approaches. In this paper, we determine these limits with a system identification procedure, similarly to that proposed in [5]. Whereas the limits were constant in time in our previous work [33], in this paper we consider time-varying bounds to better capture the impact of domestic hot water demand and EWH availability.

The VBM is then fed into the scheduling step, whose goal is to optimize the operations of the pool of appliances to provide aFRRs. Due to the fact that this optimization happens at pool level with an aggregate model, the complexity of the problem is decoupled from the number of appliances in the cluster. The scheduling framework is the focus of this paper. It consists of two building blocks, which are further detailed in the next section: the uncertainty model for aFRR activations, and the optimizer. The scheduler computes

a power baseline and aFRR bids for the pool, while ensuring that the state of charge of the virtual battery stays within limits. The bids are used by the TSO to compute aFRR activation setpoints for the pool of heaters in real time.

The bid prices of the aggregator are determined based on a pricing strategy, which can depend on different factors. Examples include the operational expenditure of the aggregator, the opportunity cost for the provision of another service, or the availability of assets. Based on certain knowledge about the MO curve, the aggregator can estimate its MO position(s) based on the price(s) it chooses to set.

Finally, the baseline and activation setpoints computed at pool level are tracked in real time in the dispatcher. The role of the dispatcher is to accurately follow the aggregate power reference by switching heaters ON and OFF, while ensuring the individual appliance constraints are respected. We adopt the same priority-stack-based dispatch as in [33], which is based on the idea presented in [36]: the EWHs are ranked according to their need to consume power in the short term, which is a function of the current water temperature, the minimum comfort temperature, and the predicted domestic hot water consumption. When the power needs to be modulated, EWHs with high priority are switched ON, while EWHs with low priority are switched OFF. Note that the dispatch does not overrule the comfort and safety rules of Equation (1). In the case where the power profile consisting of the baseline and the aFRR activations cannot be tracked without safety or comfort violations, it is assumed that the aggregator takes the responsibility for this control failure. Therefore, in such a case, the aggregator will default on its activation and face possible penalties, but the comfort of the appliance users is guaranteed.

As shown in Figure 4, the overall architecture results in two control loops: the inner real-time control loop is executed every 4 s to match the resolution of the aFRR control signal. The outer scheduling loop is executed every 15 min to match the resolution of the aFRR energy bids.

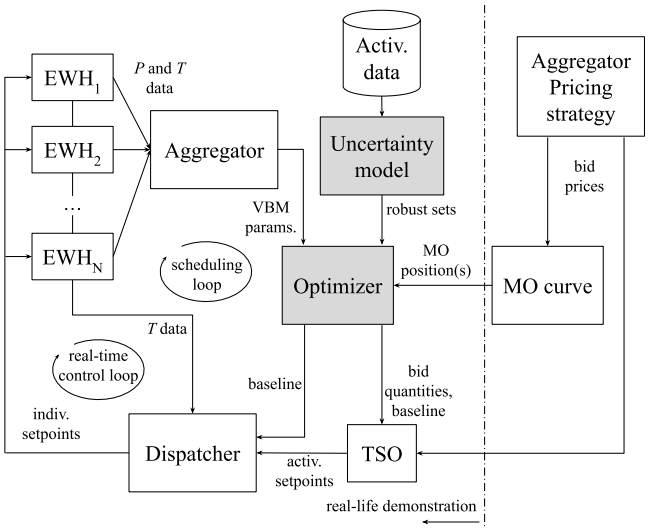

Figure 4. Control structure overview. The gray boxes show the components which are part of the scheduling step (focus of this paper).

4. Scheduling Framework

4.1. Uncertainty Model

If an aFRR energy bid with power P_{bid} is placed for a given quarter of an hour, the resulting activated energy can be seen as a random quantity which can vary from 0 to $P_{bid} \times 15$ min. As the TSO activates the energy bids according to the MO, the statistical properties of this random variable depend on the aggregator's MO position. To cope with

this, we developed an approach inspired from [29], which proposes modeling the activation separately at different MO levels. At a given MO position m for a time step k starting at time t_k, the aFRR activation energy normalized by the aggregator a_m can be computed as follows:

$$a_{m,k} = \frac{\int_{t_k}^{t_{k+1}} \max(0, \min(e(\tau) - m, P_{bid,k}))d\tau}{P_{bid,k}\Delta t} \tag{4}$$

where Δt is the length of the time interval. We want to create an uncertainty set $\mathcal{A}_m$ for the vector $\boldsymbol{a}_m$ that accurately represents the activation uncertainty in a tractable optimization problem. We build this set based on statistical deviation measures called forward and backward deviations. Such a set was originally proposed in [37] and used in [24,33] for primary reserves.

As guarantees on the confidence level of the set are only valid for independent variables with zero mean, we first transform the vector $\boldsymbol{a}_m$, as follows:

$$\widetilde{\boldsymbol{a}}_m = \boldsymbol{W}(\boldsymbol{a}_m - \bar{\boldsymbol{a}}_m) \tag{5}$$

where $\boldsymbol{W}$ is a whitening matrix computed as the Cholesky decomposition of the inverse of the covariance matrix of $\boldsymbol{a}_m - \bar{\boldsymbol{a}}_m$. An uncertainty set with confidence level ϵ can then be defined as follows:

$$\begin{aligned} \mathcal{A}_{m,\epsilon} = \{\widetilde{\boldsymbol{a}}_m : \exists \boldsymbol{\beta}, \boldsymbol{\theta} \in \mathbb{R}_+^{n_k}, \widetilde{\boldsymbol{a}}_m = \boldsymbol{\beta} - \boldsymbol{\theta}, \\ \|\boldsymbol{Q}^{-1}\boldsymbol{\beta} + \boldsymbol{R}^{-1}\boldsymbol{\theta}\| \leq \sqrt{-2\ln(\epsilon)}\} \end{aligned} \tag{6a}$$

$$\boldsymbol{Q} = \mathrm{diag}(\sigma_{f1}, ..., \sigma_{fn_k}) \tag{6b}$$

$$\boldsymbol{R} = \mathrm{diag}(\sigma_{b1}, ..., \sigma_{bn_k}) \tag{6c}$$

where σ_f and σ_b are the forward and backward deviation metrics, and n_k denotes the number of time steps. The vector $\bar{\boldsymbol{a}}_m$ and the matrices $\boldsymbol{W}$, $\boldsymbol{Q}$, and $\boldsymbol{R}$ can be estimated with historical data. Several norm definitions can be considered in (6a), resulting in different approximations of the confidence level and complexity of the set. Whereas the Euclidean norm was used in [24,33], we will use the $l_1 \cap l_\infty$ norm proposed as an alternative in [37] and defined as

$$\|\boldsymbol{v}\|_{1\cap\infty} = \max\left(\frac{1}{\sqrt{n_k}}\|\boldsymbol{v}\|_1, \|\boldsymbol{v}\|_\infty\right) \tag{7}$$

Although this norm can cause the uncertainty estimation to be more conservative, it allows us to define $\mathcal{A}_{m,\epsilon}$ as a polyhedral set, resulting in linear constraints on the optimization problems of the next subsection.

4.2. Optimization Model

The goal of the optimization problem is to maximize the aggregator's profits from the aFRR service under uncertain activation, taking into account the cost of consuming electricity for the EWH owner. Although the bids and baseline decisions need to be taken only 25 min and 1 min in advance, respectively, it is important that the optimizer is aware of their impact on the operations of the pool in the following hours. Therefore, we define

the problem for a horizon of 24 h and a time step of 15 min. In its most general form, this problem can be written as follows:

$$\min \quad \mathbb{E}[\sum_{k\in\mathcal{K}}(p_e P_k \Delta t - r_k a_k P_{bid,k}\Delta t)] \tag{8a}$$
$$\text{s.t.} \quad \boldsymbol{P} = \boldsymbol{P}_{base} - \text{diag}(\boldsymbol{P}_{bid})\boldsymbol{a} \tag{8b}$$
$$\boldsymbol{x} = \boldsymbol{d} + \boldsymbol{D}\boldsymbol{P} \tag{8c}$$
$$\boldsymbol{P}_{min} \leq \boldsymbol{P} \leq \boldsymbol{P}_{max} \tag{8d}$$
$$\boldsymbol{x}_{min} \leq \boldsymbol{x} \leq \boldsymbol{x}_{max} \tag{8e}$$

where $\boldsymbol{d}$ and $\boldsymbol{D}$ are, respectively, a vector and a matrix of parameters, both defined so that (8c) corresponds to (3a) in discretized and vectorized form. $\mathbb{E}$ is the expectation operator, p_e denotes the electricity retail price, r denotes the aFRR activation remuneration, $\mathcal{K}$ denotes the set of time steps considered in the problem, and P_{base} denotes the power baseline. In the objective (8a), the first term is the retail cost of electricity incurred by the EWHs' users, while the second term represents the aggregator's revenues from aFRR delivery, taken with the opposite sign to ensure correct maximization.

It was desirable to progressively increase the complexity of the algorithm tested in the field. Therefore, in the context of the pilot study, Problem (8a)–(8e) was reformulated in three different ways, with distinct sets of assumptions. This resulted in three different optimization problems, further referred to as Problems 1, 2, and 3 (P1-P2-P3), which are described below.

4.2.1. Problem 1

Our goal for P1 is to define a simple, deterministic problem that provides robust outputs in the field. We assume that the aggregator holds a static MO position m^*. This allows us to rewrite (8a)–(8c) based on the expected (or average) activation in that position. We also make safe assumptions, to hedge against uncertain activations: for the power limits, we assume that the bid power needs to be available at all times. For the energy limits, we make sure that the the cluster can withstand the theoretical minimum and maximum activations over the optimization horizon. This results in the following problem, in which $\overline{[\cdot]}$ denotes expected (or average) quantities:

$$\textbf{P1:} \quad \min \quad \sum_{k\in\mathcal{K}}(p_e \overline{P}_k \Delta t - r_k \overline{a}_{m^*,k} P_{bid,k}\Delta t) \tag{9a}$$
$$\text{s.t.} \quad \overline{\boldsymbol{P}} = \boldsymbol{P}_{base} - \text{diag}(\boldsymbol{P}_{bid})\overline{\boldsymbol{a}}_{m^*} \tag{9b}$$
$$\boldsymbol{P}_{base} \geq \boldsymbol{P}_{min} + \boldsymbol{P}_{bid} \tag{9c}$$
$$\boldsymbol{P}_{base} \leq \boldsymbol{P}_{max} \tag{9d}$$
$$\boldsymbol{d} + \boldsymbol{D}\boldsymbol{P}_{base} \leq \boldsymbol{x}_{max} \tag{9e}$$
$$\boldsymbol{d} + \boldsymbol{D}(\boldsymbol{P}_{base} - \boldsymbol{P}_{bid}) \geq \boldsymbol{x}_{min} \tag{9f}$$

4.2.2. Problem 2

In P2, a static MO position m^* is also assumed, but the uncertainty is modeled with the framework developed in Section 4.1, in order to obtain less conservative solutions. As we still want to guarantee the instantaneous availability of the bid power, the only constraints changed compared to P1 are the energy limits, which are formulated as chance constraints:

P2: min (9a)
s.t. (9b) − (9d)

$$\Pr(x_k \geq x_{min,k}) \geq 1 - \epsilon \quad \forall k \in \mathcal{K} \tag{10a}$$
$$\Pr(x_k \leq x_{max,k}) \geq 1 - \epsilon \quad \forall k \in \mathcal{K} \tag{10b}$$

where Pr stands for the probability. Constraints (10a) and (10b) are equivalent to the following robust constraints based on the set defined in (6a):

$$\overline{\boldsymbol{x}} = \boldsymbol{d} + \boldsymbol{D}\overline{\boldsymbol{P}} \tag{11a}$$
$$\widetilde{\boldsymbol{x}} = \overline{\boldsymbol{x}} - \boldsymbol{D}\mathrm{diag}(\boldsymbol{P}_{bid})\boldsymbol{W}^{-1}\widetilde{\boldsymbol{a}}_{m^*} \tag{11b}$$
$$\boldsymbol{x}_{min} \leq \widetilde{\boldsymbol{x}} \leq \boldsymbol{x}_{max} \qquad \forall \widetilde{\boldsymbol{a}}_{m^*} \in \mathcal{A}_{m^*,\epsilon} \tag{11c}$$

As $\mathcal{A}_{m^*,\epsilon}$ is a polyhedral uncertainty set, constraints (11b) and (11c) can be reformulated as a set of linear constraints, resulting in a tractable optimization problem.

4.2.3. Problem 3

The main goal for P3 is to discard the assumption of a static MO position. Instead, we assume that the aggregator has some knowledge about the MO and is able to link its bid price(s) to one or several MO positions and corresponding probabilities of activation. Given a discrete set of MO position options, we can formulate P3 as a generalized form of P2:

P3:

$$\min \quad \sum_{k\in\mathcal{K}}(p_e\overline{P}_k\Delta t - \sum_{m\in\mathcal{M}^*}(r_{m,k}\overline{a}_{m,k}P_{bid,m,k}\Delta t)) \tag{12a}$$
s.t. (9c) − (9d), (11a)

$$\overline{\boldsymbol{P}} = \boldsymbol{P}_{base} - \sum_{m\in\mathcal{M}^*}\mathrm{diag}(\boldsymbol{P}_{bid,m})\overline{\boldsymbol{a}}_m \tag{12b}$$
$$\widetilde{\boldsymbol{x}} = \overline{\boldsymbol{x}} - \sum_{m\in\mathcal{M}^*}\boldsymbol{D}\mathrm{diag}(\boldsymbol{P}_{bid,m})\boldsymbol{W}_m^{-1}\widetilde{\boldsymbol{a}}_m \tag{12c}$$
$$\boldsymbol{x}_{min} \leq \widetilde{\boldsymbol{x}} \leq \boldsymbol{x}_{max} \;\; \forall m \in \mathcal{M}^*, \forall \widetilde{\boldsymbol{a}}_m \in \mathcal{A}_{m,\epsilon_m} \tag{12d}$$

where $\mathcal{M}^*$ denotes the set of selected merit order position options. As the activations at different MO positions can be statistically dependent, if we want to guarantee an overall confidence level of ϵ, we need to select the confidence levels of the sets in (12d) so that $\sum_{m\in\mathcal{M}^*}\epsilon_m = \epsilon$ [38]. Therefore, we select the same level $\epsilon_m = \epsilon/n^*$ for all sets, with n^* being the cardinality of $\mathcal{M}^*$.

5. Real-Life Validation

5.1. Experimental Setup

The EWHs steered in this pilot were all existing appliances in Flemish apartment buildings and equipped with the control module described in Section 3. Table 1 summarizes the three experimental phases followed during the pilot. To prevent any uncontrolled negative impacts on the EWH users, the complexity of the algorithm and the number of EWHs involved were gradually increased. All three periods were 21 days long and had the same number of weekdays and weekend days. The numbers of EWHs in Table 1 are averages over each period. The number of heaters varied slightly over the course of the periods, e.g., due to temporary connectivity issues. The pilot took place in 2022–2023, and was carried out together with an external project partner who was responsible for the dispatching of an activation signal based on the placed bids.

The majority of EWHs involved in the pilot were subject to a lower electricity tariff at night and during weekends, and mostly consumed energy in those time intervals. Therefore, to make sure that we stayed within the prescribed limits at all times, we restricted our bidding to the same moments. The night was defined as the period between 10 p.m. and 6 a.m. The bidding limits were enforced with additional constraints in each optimization problem.

During the first and second experimental periods, a fixed MO position $m^* = 1$ MW was assumed, and the project partner selected the activation profile to be tracked in real time among a pool of model profiles. The model profiles and the selection process were designed to match the long-term probability of activation at $m^* = 1$ MW. This value placed the pool of EWHs at the beginning of the MO, which ensured frequent activations and allowed observing the pool's response over many different events. During the third period, the project partner made it possible to place bids at three different MO positions: $\mathcal{M}^* = \{1, 72, 140\}$ MW. These values were chosen as they were representative of the beginning, middle, and end of the Belgian aFRR energy MO. In this period, the activation profiles were coupled to the market needs in real time.

The activations and their uncertainty were modeled with 1 year of historical data from 2021 [34]. The optimization problem was solved with a rolling horizon every 15 min, to make optimal use of real-time information updates, while keeping decision variables fixed where needed, to respect the timings discussed in Section 2. We used an electricity price of $p_e = 0.25$ €, as this was representative of the average electricity retail price at the time of the study [39]. A confidence level $\epsilon = 10^{-2}$ was selected, as it provided a good balance between risk exposure and conservativeness.

Table 1. Experimental steps followed in this work.

Period	Opt. Model	Av. # EWHs	Av. Controlled Capacity [kW]	Av. Storage Capacity [kWh]	Purpose
13 Sept.–04 Oct.	P1	250	549	475	Establish a benchmark
15 Nov.–06 Dec.	P2	403	868	664	Validate the uncertainty model
29 Dec.–19 Jan.	P3	619	1378	1363	Validate under varying MO position

5.2. Uncertainty Model Validation

In this subsection, we validate the use of the uncertainty model presented in this paper by comparing the results of P1 and P2 according to three criteria: revenue generation, energy efficiency, and reliability.

Figure 5 shows the normalized cumulative revenues per EWH over both experimental periods. It is clear that P2 performed better than P1: the final revenues were more than 30% higher at the end of the second period. This is because, whereas P1 considered the theoretical extremes of the aFRR activations given the placed bids, P2 considered activation scenarios based on historical data, which resulted in a less conservative bidding behavior. It can also be seen that the performance difference between the two models was the largest during the longer periods for which bidding was allowed, i.e., the weekends. This behavior was expected: as the probability of obtaining very high (or very low) consecutive activations drops over longer time frames, the difference between the theoretical worst case and the data-driven uncertainty model increases.

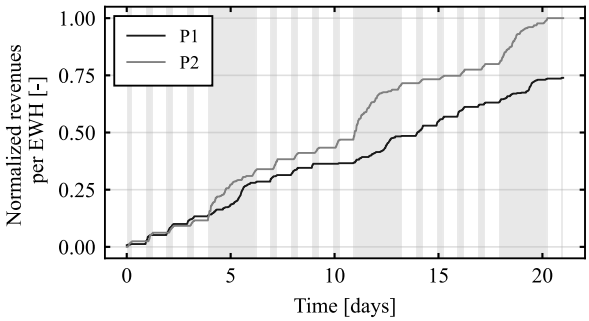

Figure 5. Cumulative normalized revenues for P1 and P2. The shaded areas show the moments when bids were allowed.

To evaluate the performance of the two approaches in terms of energy efficiency, we computed the achieved energy savings for each EWH. These savings were obtained by comparing the actual EWH consumption with a model-based simulation of what the EWH would have consumed without external control. Under its default settings, an EWH is controlled thermostatically within a deadband close to its maximum temperature. More specifically, the EWH state was determined by (1), with $T_{min,i} = T_{max,i} - 5\ ^{\circ}\mathrm{C}$.

The savings achieved across EWHs for all three periods are shown in Figure 6. Compared to P1, P2 achieved an increase in average savings of approximately 2 percentage points. This difference can be attributed to the fact that the constraints of P2 were less restrictive than those of P1. This implies that P2 had a larger decision space to minimize its objective function. The savings variations across appliances were mostly due to their usage. An EWH with a high hot water consumption compared to its storage size needs to be maintained close to its maximum temperature all the time and has a low savings potential. On the other hand, an EWH with a low hot water demand can be maintained at a much lower temperature than under its default steering, and it has a greater margin for efficiency increase.

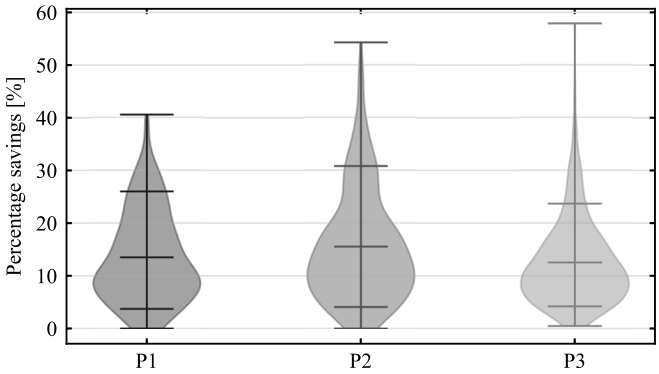

Figure 6. Violin plots of the energy savings achieved across EWHs for the different experimental periods. Horizontal lines from top to bottom: (i) maximum, (ii) 90th percentile, (iii) mean, (iv) 10th percentile, (v) minimum.

Finally, the reliability of the control in the first and second periods is summarized in Table 2. The ex-ante ϵ is the confidence level set in the optimization model. The ex-post ϵ is the ratio of the quarter of an hours during which the state of charge went out of bounds. The tracking failure percentage is the ratio of the quarter of an hours during which the pool could not accurately track its baseline or fulfill an activation. As expected, the P1 conservative approach resulted in no constraint violations or operational failures. Although constraint violations were tolerated in up to 1% of the cases for P2, the second test period did not exhibit any violations or failures either.

Table 2. Reliability analysis of periods 1 and 2.

Opt. Model	Ex-Ante ϵ	Ex-Post ϵ	Tracking Failure Percentage
P1	0%	0%	0%
P2	1%	0%	0%

5.3. Validation Under Varying Merit Order Positions

The goal of the third experimental period was to validate P3 when the bid prices set by the aggregator resulted in different MO positions at different moments. To enable this, the project partner modified the activation method to make it dependent on the MO position. However, this new method resulted in activation volumes that were significantly lower than those of the real market. Due to these altered conditions, the performance of P3 cannot be reliably compared with that of P1 and P2. Nevertheless, the test period allows us to illustrate the feasibility of P3 for a large EWH cluster, and provides interesting insights.

Figure 7 shows the operations of the EWH pool under P3 for 2 weekdays and 1 weekend day. At the end of the weekday nights, the EWHs needed to reach a high temperature to cover the hot water needs of the user during the following day. That is why the lower state of charge limit x_{min} rose towards the end of the night. The bottom plot shows the power of the individual EWHs in grayscale.

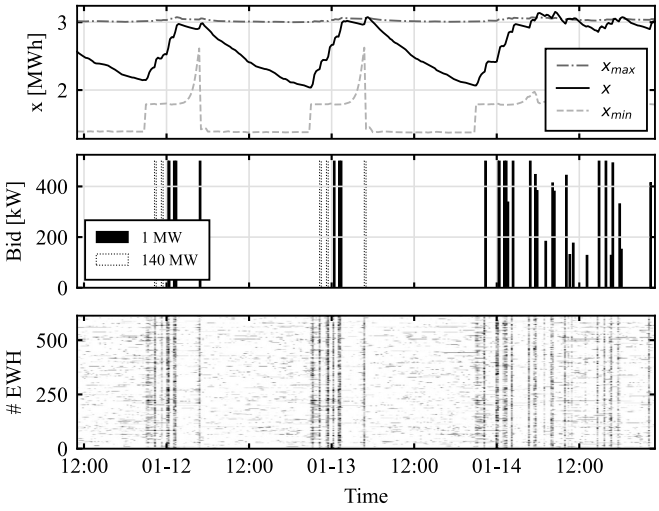

Figure 7. Operational example from January 2023 (P3). From top to bottom: (i) state of charge of the cluster with its minimum and maximum limits, (ii) bids placed at each MO position, and (iii) individual EWH consumptions.

Two consecutive successful aFRR activations can be examined in more detail in Figure 8. Based on the communicated baseline and bid, the project partner ordered the EWH cluster to follow the dashed power reference in real time. It can be seen that a very accurate tracking was achieved, which also validated the performance of the dispatcher (common to all test periods).

Due to the change in activation method, the activation volumes in the third period tended to be overestimated by our market-calibrated uncertainty model. Due to this distortion, the reliability performance of P3 could not be accurately estimated, and is therefore not included in Table 2.

This overestimation of the activation volumes probably explains the slightly lower energy efficiency results of P3 compared to P1 and P2 in Figure 6. In any case, additional tests are required for a proper performance comparison between the approaches.

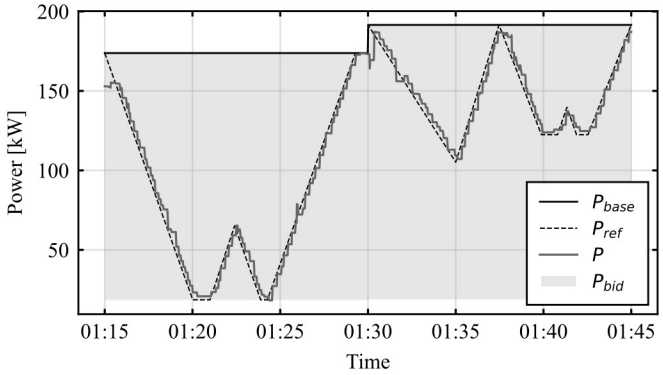

Figure 8. aFRR activation example.

6. Conclusions

In this paper, we performed the first ever real-life validation of secondary frequency reserve provision with an aggregation of thermal loads in a near-commercial setting. We proposed an aggregate-and-dispatch approach to steer a cluster of EWHs, and we focused on the development of a novel scheduling framework with two building blocks: an uncertainty model, and optimization model.

The aggregate-and-dispatch approach provided scalable control for the pool, while achieving highly accurate power reference tracking, and therefore proved adequate to tackle the control problem. The handling of the aFRR activation uncertainty with chance constraints induced a significant performance increase compared to a fully robust approach: it resulted in simultaneously higher aFRR revenues and a better energy efficiency of the EWHs, without compromising service reliability. Based on the same uncertainty model, optimization model P3 was validated as a way to allow varying positions in the MO.

The largest pool used in this work consisted of 619 EWHs. Since the control design inherently decouples scheduling complexity from the number of appliances, the presented method can be applied to much larger clusters. However, since the dispatch step involves generating individual instructions for each device, this may eventually become a bottleneck as the pool size grows. In that case, further scaling could be achieved by splitting the pool into distinct clusters that operate in parallel, which is straightforward with modern cloud infrastructure.

In summary, we demonstrated that, with the right tools, residential electric water heaters can deliver secondary frequency reserves with high performance and reliability. However, the scheduling and control aspects are not the only challenges involved in the participation of residential appliances in frequency services. The impact on the distribution grid, information exchange, and validation of activations are examples of elements that require further investigation and development. To address these challenges, future work could explore improved grid-aware control strategies to mitigate potential distribution network constraints, compare communication protocols to identify the most suitable options for distributed energy resources such as EWHs, and review existing measurement and verification frameworks to address barriers to residential aggregations. These investigations should be conducted in close collaboration with key stakeholders (aggregators, system operators, regulators, end users, etc.) to facilitate market access.

Beyond these technical and regulatory challenges, understanding the economic value of aFRR provision with EWHs is another important research avenue. While a full economic assessment is beyond the scope of this work, two key system benefits can be highlighted. Firstly, the entry of new participants into the aFRR market enhances liquidity and fosters greater competition. Secondly, EWHs can provide frequency support without inherent

CO_2 emissions, contributing to decarbonization efforts. Evaluating these economic and environmental impacts in greater detail remains an essential area for future work.

Finally, future research could also explore extending this framework to aFRR capacity auctions and combining aFRR provision with other services, such as primary frequency reserves.

Author Contributions: Conceptualization, L.B.; methodology, L.B.; software, L.B.; validation, L.B., S.I. and S.Q.; formal analysis, L.B.; investigation, L.B.; resources, L.B., S.I. and S.Q.; data curation, L.B.; writing—original draft preparation, L.B.; writing—review and editing, S.I. and S.Q.; visualization, L.B.; supervision, S.I. and S.Q.; project administration, S.I.; funding acquisition, S.I. All authors have read and agreed to the published version of the manuscript.

Funding: This research was partially funded by the Energy Transition Fund from the Belgian Federal Public Service Economy, under the grant name UNLEASH.

Data Availability Statement: The datasets presented in this article are not readily available because of commercial and privacy restrictions.

Conflicts of Interest: Author Sandro Iacovella is the chief executive officer of ThermoVault BV. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Nomenclature

$[\cdot]_{min,max}$	Minimum and maximum limits.
$\overline{[\cdot]}$	Average quantity.
a	Normalized aFRR activation by the aggregator, uncertainty set $\mathcal{A}$.
α, β	Thermal model parameters.
Δt	Time step duration.
$\mathbb{E}$	Expectation operator.
e	Instantaneous aFRR system activation.
k	Time step index, set $\mathcal{K}$, cardinality n_k.
m	Merit order position.
$\mathcal{M}^*$	Set of selected merit order position options, cardinality n^*.
n_{ewh}	Number of electric water heaters, index i.
Pr	Probability.
P, P_{base}, P_{bid}	Pool power, baseline power, bid power.
P_{ref}	Tracking power reference.
P_{rated}	Electric water heater rated power.
P_w	Equivalent power associated with domestic hot water consumption.
p_e	Electricity retail price.
p_{bid}	Bid price.
r	aFRR activation remuneration.
σ_f, σ_b	Forward and backward deviations.
s	Relay state.
T, T_a	Temperature, ambient temperature.
t	Time.
$\boldsymbol{v}$	A vector with length n_k and scalar components v_k.
$\boldsymbol{V}$	A matrix with dimensions (n_k, n_k).
x	State of charge of the virtual battery model.

References

1. Callaway, D.S.; Hiskens, I.A. Achieving Controllability of Electric Loads. *Proc. IEEE* **2011**, *99*, 184–199. [CrossRef]
2. Kemna, R.; van Elburg, M.; Aarts, S.; Corso, A. *Water Heaters Ecodesign and Energy Label—Preparatory Review Study—Task 2: Market Analysis*; Technical Report; European Commission: Brussels, Belgium, 2019.
3. Callaway, D.S. Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy. *Energy Convers. Manag.* **2009**, *50*, 1389–1400.

4. Koch, S.; Mathieu, J.L.; Callaway, D.S. Modeling and Control of Aggregated Heterogeneous Thermostatically Controlled Loads for Ancillary Services. In Proceedings of the 17th Power Systems Computation Conference (PSCC'11), Stockholm, Sweden, 22–26 August 2011.
5. Mathieu, J.L.; Kamgarpour, M.; Lygeros, J.; Andersson, G.; Callaway, D.S. Arbitraging Intraday Wholesale Energy Market Prices With Aggregations of Thermostatic Loads. *IEEE Trans. Power Syst.* **2015**, *30*, 763–772.
6. Iacovella, S.; Ruelens, F.; Vingerhoets, P.; Claessens, B.; Deconinck, G. Cluster Control of Heterogeneous Thermostatically Controlled Loads Using Tracer Devices. *IEEE Trans. Smart Grid* **2017**, *8*, 528–536. [CrossRef]
7. Vandael, S.; Claessens, B.; Hommelberg, M.; Holvoet, T.; Deconinck, G. A Scalable Three-Step Approach for Demand Side Management of Plug-in Hybrid Vehicles. *IEEE Trans. Smart Grid* **2013**, *4*, 720–728.
8. Biegel, B.; Andersen, P.; Stoustrup, J.; Madsen, M.B.; Hansen, L.H.; Rasmussen, L.H. Aggregation and Control of Flexible Consumers—A Real Life Demonstration. *IFAC Proc. Vol.* **2014**, *47*, 9950–9955.
9. Wrinch, M.; Dennis, G.; El-fouly, T.H.M.; Wong, S. Demand Response Implementation for Improved System Efficiency in Remote Communities Pilot Results from the Village of Hartley Bay. In Proceedings of the IEEE annual Electrical Power and Energy Conference (EPEC 2012), London, ON, Canada, 10–12 October 2012; pp. 105–110.
10. Müller, F.; Jansen, B. Large-scale demonstration of precise demand response provided by residential heat pumps. *Appl. Energy* **2019**, *239*, 836–845.
11. Vrettos, E.; Kara, E.C.; MacDonald, J.; Andersson, G.; Callaway, D.S. Experimental Demonstration of Frequency Regulation by Commercial Buildings—Part I: Modeling and Hierarchical Control Design. *IEEE Trans. Smart Grid* **2018**, *9*, 3213–3223.
12. Vrettos, E.; Kara, E.C.; MacDonald, J.; Andersson, G.; Callaway, D.S. Experimental Demonstration of Frequency Regulation by Commercial Buildings—Part II: Results and Performance Evaluation. *IEEE Trans. Smart Grid* **2018**, *9*, 3224–3234.
13. Qureshi, F.A.; Lymperopoulos, I.; Khatir, A.A.; Jones, C.N. Economic Advantages of Office Buildings Providing Ancillary Services with Intraday Participation. *IEEE Trans. Smart Grid* **2018**, *9*, 3443–3452. [CrossRef]
14. Iria, J.; Soares, F.; Matos, M. Optimal bidding strategy for an aggregator of prosumers in energy and secondary reserve markets. *Appl. Energy* **2019**, *238*, 1361–1372.
15. Vrettos, E.; Oldewurtel, F.; Andersson, G. Robust Energy-Constrained Frequency Reserves From Aggregations of Commercial Buildings. *IEEE Trans. Power Syst.* **2016**, *31*, 4272–4285. [CrossRef]
16. Almasalma, H.; Deconinck, G. Simultaneous Provision of Voltage and Frequency Control by PV-Battery Systems. *IEEE Access* **2020**, *8*, 152820–152836. [CrossRef]
17. Margellos, K.; Goulart, P.; Lygeros, J. On the road between robust optimization and the scenario approach for chance constrained optimization problems. *IEEE Trans. Autom. Control* **2014**, *59*, 2258–2263.
18. Karasavvidis, M.; Stratis, A.; Papadaskalopoulos, D.; Strbac, G. Optimal Offering of Energy Storage in UK Day-ahead Energy and Frequency Response Markets. *J. Mod. Power Syst. Clean Energy* **2024**, *12*, 415–426.
19. Herre, L.; Soder, L.; Mathieu, J.L. The Flexibility of Thermostatically Controlled Loads as a Function of Price Notice Time. In Proceedings of the 2018 Power Systems Computation Conference (PSCC), Dublin, Ireland, 11–15 June 2018.
20. Namor, E.; Sossan, F.; Cherkaoui, R.; Paolone, M. Control of Battery Storage Systems for the Simultaneous Provision of Multiple Services. *IEEE Trans. Smart Grid* **2019**, *10*, 2799–2808. [CrossRef]
21. Amini, M.; Almassalkhi, M. Optimal Corrective Dispatch of Uncertain Virtual Energy Storage Systems. *IEEE Trans. Smart Grid* **2020**, *11*, 4155–4166.
22. Almasalma, H.; Deconinck, G. Robust Policy-Based Distributed Voltage Control Provided by PV-Battery Inverters. *IEEE Access* **2020**, *8*, 124939–124948. [CrossRef]
23. Zhang, H.; Hu, Z.; Munsing, E.; Moura, S.J.; Song, Y. Data-Driven Chance-Constrained Regulation Capacity Offering for Distributed Energy Resources. *IEEE Trans. Smart Grid* **2019**, *10*, 2713–2725.
24. Engels, J.; Claessens, B.; Deconinck, G. Optimal Combination of Frequency Control and Peak Shaving with Battery Storage Systems. *IEEE Trans. Smart Grid* **2020**, *11*, 3270–3279.
25. Poolla, B.K.; Hota, A.R.; Bolognani, S.; Callaway, D.S.; Cherukuri, A. Wasserstein Distributionally Robust Look-Ahead Economic Dispatch. *IEEE Trans. Power Syst.* **2021**, *36*, 2010–2022. [CrossRef]
26. Chen, W.; Sim, M. Goal-Driven Optimization. *Oper. Res.* **2009**, *57*, 342–357. [CrossRef]
27. The European Commission. Commission Regulation (EU) 2017/2195 of 23 November 2017 establishing a guideline on electricity balancing. *Off. J. Eur. Union* **2017**, *L 312*, 6–53.
28. Schillemans, A.; De Vivero Serrano, G.; Bruninx, K. Strategic Participation of Merchant Energy Storage in Joint Energy-Reserve and Balancing Markets. In Proceedings of the Mediterranean Conference on Power Generation, Transmission, Distribution and Energy Conversion (MEDPOWER 2018), Dubrovnik, Croatia, 12–15 November 2018.
29. Toubeau, J.F.; Bottieau, J.; De Greeve, Z.; Vallee, F.; Bruninx, K. Data-Driven Scheduling of Energy Storage in Day-Ahead Energy and Reserve Markets with Probabilistic Guarantees on Real-Time Delivery. *IEEE Trans. Power Syst.* **2021**, *36*, 2815–2828. [CrossRef]

30. Dolanyi, M.; Bruninx, K.; Toubeau, J.F.; Delarue, E. Capturing Electricity Market Dynamics in Strategic Market Participation using Neural Network Constrained Optimization. *IEEE Trans. Power Syst.* **2023**, *39*, 533–545.
31. Aneiros, G.; Vilar, J.M.; Cao, R.; Munõz San Roque, A. Functional prediction for the residual demand in electricity spot markets. *IEEE Trans. Power Syst.* **2013**, *28*, 4201–4208.
32. Campos, F.A.; Munoz San Roque, A.; Sanchez-Ubeda, E.F.; Portela Gonzalez, J. Strategic Bidding in Secondary Reserve Markets. *IEEE Trans. Power Syst.* **2016**, *31*, 2847–2856. [CrossRef]
33. Brouyaux, L.; Iacovella, S.; Olivella-Rosell, P.; Quoilin, S. Chance-Constrained Frequency Containment Reserves Scheduling with Electric Water Heaters. In Proceedings of the 2021 IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe), Espoo, Finland, 18–21 October 2021.
34. Elia. Grid Data—Balancing. Available online: https://www.elia.be/en/grid-data/balancing (accessed on 10 March 2025).
35. Brouyaux, L.; Olivella-Rosell, P.; Iacovella, S.; Quoilin, S. Comparative analysis of aggregate battery models to characterize the flexibility of electric water heaters. In Proceedings of the ECOS 2021—The 34th International Conference on Efficiency, Cost, Optimization, Simulation and Environmental Impact of Energy Systems, Taormina, Italy, 27 June–2 July 2021.
36. Hao, H.; Sanandaji, B.M.; Poolla, K.; Vincent, T.L. Aggregate Flexibility of Thermostatically Controlled Loads. *IEEE Trans. Power Syst.* **2015**, *30*, 189–198. [CrossRef]
37. Chen, X.; Sim, M.; Sun, P. A Robust Optimization Perspective on Stochastic Programming. *Oper. Res.* **2007**, *55*, 1058–1071. [CrossRef]
38. Bertsimas, D.; Gupta, V.; Kallus, N. Data-driven robust optimization. *Math. Program.* **2018**, *167*, 235–292. [CrossRef]
39. VREG. Hoeveel Betaalt u Also u Vandaag Been Nieuw Commercieel Elektriciteitscontract Afsluit? Available online: https://dashboard.vreg.be/report/DMR_Prijzen_elektriciteit.html (accessed on 10 March 2025).