

Reliable Frequency Regulation through Vehicle-to-Grid: From EU Legislation to Robust Optimization

Dirk Lauinger^{*1}, François Vuille¹, and Daniel Kuhn¹

¹*Ecole polytechnique fédérale de Lausanne, Switzerland*

September 3, 2021

Abstract

Vehicle-to-grid increases the low utilization rate of privately owned electric vehicles by making their batteries available to the electricity grid. We formulate a robust optimization problem with functional uncertainties that maximizes the expected profit from selling primary frequency regulation to the grid and guarantees that vehicle owners can meet their market commitments for all frequency deviation trajectories in an uncertainty set that encodes applicable EU legislation. Faithfully modeling the energy conversion losses during battery charging and discharging renders this optimization problem non-convex. By exploiting a total unimodularity property of the proposed uncertainty sets and an exact linear decision rule reformulation, we prove that this non-convex robust optimization problem with functional uncertainties is equivalent to a tractable linear program. Through extensive numerical experiments based on real-world data, we faithfully quantify the value of vehicle-to-grid, which helps us to understand the economic incentives of vehicle owners, aggregators, equipment manufacturers and regulators.

Keywords: Vehicle-to-Grid, Frequency Regulation, Energy Storage, Energy Economics, Robust Optimization, Continuous-Time Linear Programming.

1 Introduction

Replacing internal combustion engine vehicles with electric vehicles reduces urban air pollution and mitigates climate change if electricity is generated from renewable sources (Sperling, 1994). In

^{*}Corresponding author: dirk.lauinger@epfl.ch

general, privately owned vehicles are a vastly underutilized resource. Vehicle usage data collected by the US Federal Highway Administration (2017) shows that on an average day over 90% of all privately owned vehicles are parked at any one time—even during peak rush hour. Since electricity grids require storage capacity to integrate increasing amounts of intermittent wind and solar power, electric vehicle owners could capitalize on their batteries by offering storage to the electricity grid when their vehicles are parked. Kempton and Letendre (1997) term this idea *vehicle-to-grid*.

Réseau de transport d'électricité (RTE), Europe's largest transmission system operator, expects to need an additional flexible generation and electricity storage capacity of 10GW to 20GW by 2035. This corresponds to 7.5% to 15% of the total French electricity generation capacity in 2017 (RTE 2017b; 2017c). If electric vehicles were to provide some of this flexibility, then the vehicles and the electricity grid could share the costs of electric vehicle batteries. Kempton and Tomić (2005) and Noel et al. (2019) have identified *primary frequency regulation*¹ as one of the most profitable flexibility services for vehicle-to-grid. Electric vehicles that provide this service must maintain a continuous power flow to the vehicle battery that is proportional to the deviation of the instantaneous grid frequency from its nominal value (*i.e.*, 50Hz in Europe). As primary frequency regulation is the first flexibility service used to stabilize the electricity network after disturbances (Rebours et al., 2007), its provision must be highly reliable. However, RTE questions the reliability of vehicle-to-grid (RTE 2017b). The European Commission (2017) has recently addressed this concern by defining a minimum level of reliability that electric vehicles and other providers of frequency regulation must guarantee. Specifically, it demands that providers must be able to deliver regulation power for all frequency deviation trajectories with certain characteristics.

Adopting the perspective of a vehicle owner, we formulate an optimization model for determining the bidding strategy on the regulation market that maximizes the expected profit from selling primary frequency regulation to the transmission system operator under the reliability constraints imposed by the European Commission. These constraints must hold *robustly* for all frequency deviation trajectories in an uncertainty set consistent with applicable legislation. As these trajectories constitute continuous-time functions, we are confronted with a robust optimization problem with functional uncertainties. Moreover, the impossibility of simultaneously charging and discharging the battery—which amounts to dissipating energy through conversion losses and could be profitable

¹Primary frequency regulation is also known as primary frequency control and as frequency containment reserves.

when the battery is full and there is a reward for down-regulation (see, *e.g.*, (Taylor, 2015, p. 84))—renders the optimization problem non-convex. The main theoretical contribution of this paper is to show that the resulting non-convex robust optimization problem with functional uncertainties is equivalent to a tractable linear program. Specifically, this paper makes the following methodological contributions to robust optimization (see Ben-Tal et al. (2009) for a textbook introduction).

- We introduce new *uncertainty sets in function spaces* that capture those frequency deviation trajectories for which regulation providers must be able to deliver all promised regulation power. These uncertainty sets are reminiscent of the *budget uncertainty sets* by Bertsimas and Sim (2004) in finite-dimensional spaces, and their construction is inspired by EU legislation.
- By leveraging a *total unimodularity property* of the proposed uncertainty sets and an *exact linear decision rule reformulation*, we prove that the worst-case frequency deviation scenarios in all (convex or non-convex) robust constraints of the vehicle owner’s optimization problem can be found by solving continuous linear programs, which can be viewed as variants of the so-called *separated continuous linear programs* introduced by Anderson et al. (1983).
- By demonstrating that all these continuous linear programs are solved by piecewise constant frequency deviation trajectories, we show that the vehicle owner’s robust optimization problem in continuous time is equivalent to a *robust optimization problem in discrete time*. In doing so, we use more direct proof techniques than Pullan (1995), who derived sufficient conditions under which the solutions of separated continuous linear programs are piecewise constant.
- The robust optimization problem obtained by time discretization is still non-convex. Using the structural properties of its (discretized) uncertainty sets and of its objective and constraint functions, however, we can prove that it is equivalent to a *linear* robust optimization problem that can be reformulated as a *tractable linear program* via standard techniques.

To our best knowledge, robust optimization models with uncertainty sets embedded in function spaces have so far only been considered in the context of robust control, where the primary goal is to develop algorithms for evaluating conservative approximations (Houska, 2011), and in the context of robust continuous linear programming, where the primary goal is to reduce robust to *non*-robust continuous linear programs, which can be addressed with existing algorithms (Ghate, 2020). In contrast, we study here a non-convex robust optimization problem with functional uncertainties

that admits a lossless time discretization and can be reformulated *exactly* as a tractable linear program. Remarkably, the state-of-the-art methods for solving the deterministic counterparts of this robust optimization problem are based on methods from mixed-integer linear programming. To our best knowledge, we thus describe the first class of practically relevant mixed-integer linear programs that simplify to standard linear programs through robustification.

As the emerging linear programs are amenable to efficient numerical solution, we are able to perform extensive numerical experiments based on real-world data pertaining to the French electricity system. We define the *value of vehicle-to-grid* as the profit from selling primary frequency regulation relative to a baseline scenario in which the vehicle owner does not offer grid services. As our optimization model faithfully captures effective legislation, it enables us to quantify the true value of vehicle-to-grid. This capability is relevant for understanding the economic incentives of different stakeholders such as vehicle owners, aggregators, equipment manufacturers, and regulators. The model developed in this paper enables us to assess how the value of vehicle-to-grid depends on the penalties for non-delivery of promised regulation power, the size of the uncertainty set, and the vehicle’s battery, charger, and mileage. We thus contribute to the growing literature on the impact of contract parameters on electricity storage (Broneske and Wozabal, 2017; Sunar and Birge, 2019). The main insights drawn from our computational experiments can be summarized as follows.

- Based on 2016–2019 data, we show that the value of vehicle-to-grid attainable with a bidding strategy that is *guaranteed* to satisfy all reliability requirements is around 100€ per year and vehicle. Earlier studies based on anticipative bidding strategies that may violate the legal requirements in practice have estimated this value to be four times higher (Codani et al., 2015; Borne, 2019).
- We find a similar value of vehicle-to-grid as Codani et al. (2015) and Borne (2019) if the vehicle owner risks financial penalties for ignoring the legal reliability requirements. This suggests that *current penalties are too low* to incentivize vehicle owners to respect the law.
- We show that the value of vehicle-to-grid saturates at daily plug times above 15 hours. Thus, maximal profits from frequency regulation can be reaped even if the vehicle is disconnected from the grid up to 9 hours per day. This means that vehicle owners still enjoy considerable flexibility as to when to drive, which could help to promote the adoption of vehicle-to-grid.

Beyond vehicle-to-grid, this paper contributes to the literature on the optimal usage of energy storage assets. The value of a storage asset is usually identified with the profit that can be generated through arbitrage by trading the stored commodity on spot or forward markets. If trading is restricted to the spot market and prices are Markovian, it is known that the asset's value is maximized by a basestock policy (Secomandi, 2010). If the commodity is also traded on forward markets, then the high dimensional models of forward curve evolution lead to intractable Markov decision processes that can be addressed with approximate dynamic programming methods (Nadarajah et al., 2015). For systems of interconnected storage assets with large capacities such as hydroelectric reservoirs, medium-term planning over several months or years is necessary. The resulting optimization problems are traditionally addressed with stochastic dynamic programming (Yeh, 1985) or stochastic dual dynamic programming (Pereira and Pinto, 1991). Alternatively, Pritchard et al. (2005) use a two-layer dynamic programming method to optimize the participation of a hydroelectric reservoir in a spot market, where the inner layer maximizes the expected revenues over a stage, which comprises several trading intervals, for a fixed mean and variance of water release over the stage, while the outer layer optimizes the mean and variance of water release over the stages. More recently, Löhndorf et al. (2013) combine ideas from stochastic dual dynamic programming and approximate dynamic programming for optimizing the forward trading decisions of hydro storage systems.

Unlike traditional centralized storage assets, decentralized storage assets such as electric vehicles are usually connected to distribution rather than transmission grids. This means that they face retail and not wholesale electricity prices. While wholesale prices are determined by market mechanisms and thus stochastic, retail prices are often regulated and thus deterministic. Another major difference is that it may take several days to fully charge or discharge centralized storage assets such as hydropower plants, whereas the batteries of electric vehicles can be fully charged and discharged in just a few hours. A daily planning horizon is therefore sufficient for optimizing their usage. In addition, typical vehicle owners can anticipate their driving needs at most one day in advance. One can thus solve the storage management problem in a receding horizon fashion.

The state-of-charge of a vehicle battery depends non-linearly on the power in- and outflows, which leads to non-convex optimization models. If the battery is merely used for arbitrage and market prices are non-negative, then these optimization models admit exact convex relaxations. Conversely, if the battery is used for frequency regulation or if market prices can fall below zero, then

a non-convex constraint is needed to prevent the models from dissipating energy by simultaneously charging and discharging the battery (Zhou et al., 2016). If energy conversion losses are negligible and the battery state-of-charge is thus linear in the power flows, then one can model the provision of frequency regulation through adjustable uncertainty sets. Such an approach has been proposed by Zhang et al. (2017) for frequency regulation with building appliances. A stochastic dynamic programming scheme for optimizing the charging and discharging policy of an electric vehicle with linear battery dynamics is proposed by Donadee and Ilić (2014). If energy conversion losses are significant, however, one may still approximate the state-of-charge by a linear decision rule of the uncertain frequency deviations (Warrington et al., 2013). Sortomme and El-Sharkawi (2012) study a similar model under the assumption of perfect foresight.

In practice, several hundreds or thousands of electric vehicles must be aggregated to be able to bid enough reserve power to qualify for participation in the frequency regulation market. Guille and Gross (2009), Han et al. (2010), Wenzel et al. (2018) and Zhang et al. (2021) develop frameworks for controlling the batteries of aggregated vehicles, while the design of contracts between aggregators and vehicle owners is examined by Han et al. (2011) and Broneske and Wozabal (2017). The policy implications for the market entry of electric vehicle aggregators are investigated by Borne et al. (2018). Yet the study of vehicle-to-grid schemes for individual vehicles remains relevant because they constitute important building blocks for aggregation schemes and because they still pose many challenges—especially when it comes to faithfully modeling all major sources of uncertainty.

The model developed in this paper is most closely related to the discrete-time robust optimization models by Yao et al. (2017) and Namor et al. (2019), which capture the uncertainty of the frequency deviations through simplicial uncertainty sets that cover all empirical frequency deviation scenarios. However, these uncertainty sets may fail to include unseen future frequency deviation scenarios and are inconsistent with applicable EU legislation. While Yao et al. (2017) disregard energy conversion losses, Namor et al. (2019) account for them heuristically and test the resulting charging and discharging policies experimentally on a real battery. Heuristics are also common in pilot projects that demonstrate the use of vehicle-to-grid for frequency regulation (Vandael et al., 2013, 2020).

The model proposed in this paper relies on three simplifying assumptions that we justify below.

Our first key assumption is that the provision of frequency regulation has no negative impact on battery lifetime—even though the fear of battery degradation has been identified as a major obstacle

to the widespread adoption of vehicle-to-grid (Lauinger et al., 2017). To justify this assumption, we point out that the impact of vehicle-to-grid on battery longevity is not yet well understood. In fact, Dubarry et al. (2017) claim that such degradation is severe, while Uddin et al. (2017) claim that vehicle-to-grid may actually extend battery lifetime. In (Uddin et al., 2018), the authors of these two studies reconcile their contradictory findings by concluding that the impact of vehicle-to-grid depends on the operating conditions of the battery, such as its temperature and variations in its state-of-charge. We further justify our no-degradation assumption by restricting the battery state-of-charge to lie within 20% and 80% of the nominal battery capacity. Thompson (2018) suggests these restrictions as a rule of thumb for extending the lifetime of common lithium-ion batteries, and Sweda et al. (2017) adopt similar rules to optimize recharging policies of electric vehicles. Models that account for battery degradation are studied by He et al. (2016) and Carpentier et al. (2019).

Our second key assumption is that vehicle owners can specify time and energy windows for their driving needs one day in advance. This assumption makes sense for commuters who adhere to predictable daily routines, for example.

The third key assumption is that the vehicle owners are price takers who influence neither the market prices nor the grid frequency. This assumption is reasonable because one vehicle may cover at most several kilowatts of the 700 megawatts required for frequency regulation in France. A model of a regulation provider influencing the grid frequency is described by Mercier et al. (2009).

The paper proceeds as follows. Section 2 formulates the vehicle owner’s decision problem for a single day as a non-convex robust program with functional uncertainties. In Sections 3 and 4 we show that this problem can be reformulated equivalently as a non-convex robust program with vectorial uncertainties and even as a tractable linear program, respectively. Section 5 formulates a decision problem that looks several days into the future and shows that the resulting multistage model is still equivalent to a linear program. Numerical experiments are discussed in Section 6, and policy insights are distilled in Section 7. All proofs are relegated to the appendix.

Notation. All random variables are designated by tilde signs. Their realizations are denoted by the same symbols without tildes. Vectors and matrices are denoted by lowercase and uppercase boldface letters, respectively. For any $z \in \mathbb{R}$, we define $[z]^+ = \max\{z, 0\}$ and $[z]^- = \max\{-z, 0\}$ such that $z = z^+ - z^-$. The intersection of a set $\mathcal{A} \subseteq \mathbb{R}^d$ with \mathbb{R}_+^d is denoted by \mathcal{A}^+ . For any closed intervals $\mathcal{T}, \mathcal{U} \subseteq \mathbb{R}$, we define $\mathcal{L}(\mathcal{T}, \mathcal{U})$ as the space of all Riemann integrable functions $f : \mathcal{T} \rightarrow \mathcal{U}$,

and we denote the intersection of a set $\mathcal{B} \subseteq \mathcal{L}(\mathcal{T}, \mathbb{R})$ with $\mathcal{L}(\mathcal{T}, \mathbb{R}_+)$ as \mathcal{B}^+ .

2 Problem Description

Consider an electric vehicle whose state at any time t is characterized by the amount of energy $y(t)$ stored in its battery and the instantaneous power consumption for driving $d(t)$. We require that $y(t)$ is never smaller than \underline{y} and never larger than \bar{y} . To mitigate battery degradation, we set these limits to 20% and 80% of the nominal battery capacity, respectively. The battery interacts with the power grid through a bidirectional charger with charging efficiency $\eta^+ \in (0, 1]$ and discharging efficiency $\eta^- \in (0, 1]$, where an efficiency of 1 corresponds to a lossless energy conversion between the grid and the battery. The charger is further characterized by its maximum power consumption $\bar{y}^+(t)$ from the grid and its maximum power provision to the grid $\bar{y}^-(t)$. The power the battery can charge or discharge is therefore limited by $\eta^+ \bar{y}^+(t)$ and $\frac{1}{\eta^-} \bar{y}^-(t)$, respectively. Note that $\bar{y}^+(t)$ and $\bar{y}^-(t)$ depend on the charger to which the vehicle is connected at time t . When the vehicle is not connected to any charger, *e.g.*, when it is driving, then both $\bar{y}^+(t)$ and $\bar{y}^-(t)$ must vanish. A stationary battery can be modeled by setting $d(t) = 0$ and keeping $\bar{y}^+(t)$ and $\bar{y}^-(t)$ constant for all t .

In order to charge the battery at time t , the vehicle owner may buy power $x^b(t)$ from the local utility at a known time-varying price $p^b(t)$ as is the case under dynamic pricing schemes or day/night tariffs. In addition, she may also use the vehicle battery to earn extra revenue by providing primary frequency regulation, which can be viewed as an insurance bought by the transmission system operator (TSO) to balance unforeseen mismatches of electricity demand and supply in real time (Glover et al., 2010). If there is more supply than demand, the frequency of the power grid rises. Conversely, if there is more demand than supply, the frequency falls. A battery owner offering regulation power $x^r(t)$ at time t is obliged to change her nominal power consumption $x^b(t)$ from the grid by $\delta(t)x^r(t)$, where $\delta(t)$ quantifies the normalized deviation of the instantaneous grid frequency $f(t)$ from its nominal value f_0 (Gestionnaire du Réseau de Transport d'Electricité, 2009). Formally, we have

$$\delta(t) = \begin{cases} +1 & \text{if } f(t) > f_0 + \Delta f, \\ \frac{f(t)-f_0}{\Delta f} & \text{if } f_0 - \Delta f \leq f(t) \leq f_0 + \Delta f, \\ -1 & \text{if } f(t) < f_0 - \Delta f, \end{cases}$$

where $\Delta f > 0$ is a threshold beyond which all promised regulation power must be delivered.

The TSO contracts frequency regulation as an insurance over a prescribed planning horizon of length T , *e.g.*, one day. The planning horizon is subdivided into trading intervals $\mathcal{T}_k = [(k-1)\Delta t, k\Delta t)$ for all $k \in \mathcal{K} = \{1, \dots, K\}$, where $K = \frac{T}{\Delta t} \in \mathbb{N}$. In the French electricity market, for example, the length Δt of a trading interval is 30 minutes. The TSO requests the vehicle owner to announce the market decisions $x^b(t)$ and $x^r(t)$ before the beginning of the planning horizon, *e.g.*, one day ahead at noon (Réseau de transport d'électricité, 2017a). These decisions need to be piecewise constant over the trading intervals. The TSO compensates the vehicle owner for the frequency regulation $x^r(t)$ made available at the *availability price* $p^a(t)$ and charges her for the increase $\delta(t)x^r(t)$ in her power consumption at the *delivery price* $p^d(t)$. Note that this charge becomes negative (*i.e.*, it becomes a remuneration) if $\delta(t)$ is negative. In summary, the vehicle owner's total cost over the planning horizon $\mathcal{T} = [0, T]$ amounts to

$$\int_0^T p^b(t)x^b(t) - \left(p^a(t) - \delta(t)p^d(t)\right) x^r(t) dt.$$

The impact of providing frequency regulation on the battery state-of-charge depends on how the vehicle owner adjusts the power consumed from and the power injected into the grid to achieve the desired net power consumption $x^b(t) + \delta(t)x^r(t)$. The most energy-efficient way is to avoid unnecessary energy conversion losses resulting from simultaneously charging and discharging. Sometimes, however, such losses can be attractive, for example if the battery is almost full and receives a request for down-regulation ($\delta(t) > 0$). Zhou et al. (2016) show that energy losses can also be attractive when electricity prices are negative. Since common chargers are not able to simultaneously charge and discharge, we forbid this option and set the charging rate to

$$y^+ \left(x^b(t), x^r(t), \delta(t)\right) = \left[x^b(t) + \delta(t)x^r(t)\right]^+ \quad (1a)$$

and the discharging rate to

$$y^- \left(x^b(t), x^r(t), \delta(t)\right) = \left[x^b(t) + \delta(t)x^r(t)\right]^-. \quad (1b)$$

Remark 1. When operating a vehicle fleet, some vehicles could charge while others discharge, which suggests that the regulation profits achievable with n vehicles may exceed the regulation profit of a single vehicle multiplied by n . In this paper, we focus on the case $n = 1$. \square

The power exchanged with the grid and the power needed for driving determine the battery

state-of-charge at any time t via the integral equation

$$y(x^b, x^r, \delta, y_0, t) = y_0 + \int_0^t \eta^+ y^+(x^b(t'), x^r(t'), \delta(t')) - \frac{y^-(x^b(t'), x^r(t'), \delta(t'))}{\eta^-} - d(t') dt', \quad (2)$$

where y_0 represents the state-of-charge at time 0. For later use, we establish here some basic properties of the battery state-of-charge.

Proposition 1. *Holding all other factors fixed, the battery state-of-charge $y(x^b, x^r, \delta, y_0, t)$ is concave nondecreasing in x^b , concave in x^r , concave nondecreasing in δ , and linear nondecreasing in y_0 .*

At the time when the vehicle owner needs to choose and report the market commitments $x^b(t)$ and $x^r(t)$, she has no knowledge of the uncertain future frequency deviations $\delta(t)$ and the delivery prices $p^d(t)$ at time $t \in \mathcal{T}$. In addition, she has no means to predict the battery state-of-charge y_0 at the beginning of the planning horizon, which depends on market commitments chosen on the previous day and on the uncertain frequency deviations to be revealed until time 0. By contrast, the availability prices $p^a(t)$ for $t \in \mathcal{T}$ can be assumed to be known at the planning time. In practice, these prices are determined by an auction. As the vehicle owner bids an offer curve expressing $x^r(t)$ as a function of $p^a(t)$ for any $t \in \mathcal{T}$, it is as if the availability prices were known upfront.² Next, we describe the information that is available about the uncertain problem parameters δ , p^d , and y_0 .

We first discuss the uncertainty in the frequency deviations, which limits the amount of reserve power that can be sold on the market. Indeed, the vehicle owner must ensure that the battery state-of-charge will never drop below \underline{y} or exceed \bar{y} when the TSO requests down-regulation ($\delta(t) < 0$) or up-regulation ($\delta(t) > 0$), respectively, for a prescribed set of conceivable frequency deviation scenarios. Otherwise, the vehicle owner may not be able to honor her market commitments, in which case the TSO may charge a penalty or even ban her from the market.

The TSO defines under what conditions regulation providers must be able to deliver the promised regulation power, keeping in mind that extreme frequency deviations are uncommon. Indeed, between 2015 and 2018 the frequency deviation $\delta(t)$ has never attained its theoretical maximum of 1 or its theoretical minimum of -1 in the French market.³ In the following, we thus assume that the vehicle owner needs to guarantee the delivery of regulation power *only* for frequency deviation

²The bidding process is described at https://www.entsoe.eu/network_codes/eb/fcr/.

³The French TSO publicizes frequency measurements at <http://clients.rte-france.com/>.

scenarios within the uncertainty set

$$\mathcal{D} = \left\{ \delta \in \mathcal{L}(\mathcal{T}, [-1, 1]) : \int_{[t-\Gamma]^+}^t |\delta(t')| dt' \leq \gamma \quad \forall t \in \mathcal{T} \right\}$$

parametrized by the duration $\Gamma \in \mathbb{R}_+$ of a *regulation cycle* and the duration $\gamma \in \mathbb{R}_+$ of an *activation period*. Throughout this paper, we assume that $0 < \gamma \leq \Gamma \leq T$. By focusing on frequency deviation scenarios in \mathcal{D} , one stipulates that consecutive extreme frequency deviations $\delta(t) \in \{-1, 1\}$ can occur at most over one activation period within each regulation cycle. The *activation ratio* γ/Γ can thus be interpreted as the percentage of time during which the vehicle owner must be able to deliver all committed reserve power.

Remark 2. Note that the uncertainty set \mathcal{D} grows with γ and shrinks with Γ . □

Besides displaying favorable computational properties, the uncertainty set \mathcal{D} has conceptual appeal because it formalizes the delivery guarantee rules prescribed by the European Commission (2017). These rules stipulate that the “*minimum activation period to be ensured by [frequency regulation] providers [is not to be] greater than 30 or smaller than 15 minutes.*” This guideline prompts us to set $\gamma = 30$ minutes. The EU further demands that regulation providers “*shall ensure the recovery of [their] energy reservoirs as soon as possible, within 2 hours after the end of the alert state.*” This means that, although there may be several activation periods of 30 minutes within any 2.5 hour interval, the regulation provider only has to cover one of them. Thus, we set $\Gamma = 2.5$ hours.

In the following, we compare the empirical distribution of the daily variance of δ between the years 2017 and 2019 with the maximum variance that can be achieved by any hypothetical frequency deviation scenario $\delta \in \mathcal{D}$ for a planning horizon of one day. By slight abuse of notation, we define the variance of a frequency deviation scenario δ with respect to zero as $\text{Var}(\delta) = \frac{1}{T} \int_0^T \delta(t)^2 dt$. This is justified because the TSO protects the system against *unforeseen* demand and supply fluctuations, which means that the frequency deviations should be unbiased and thus vanish on average. Indeed, the empirical frequency deviations have an average of $5.98 \cdot 10^{-4}$.

Proposition 2. *If $\delta \in \mathcal{D}$, then $\text{Var}(\delta) \leq \lceil T/\Gamma \rceil \gamma/T$.*

Figure 1 shows that if $T = 1$ day, $\gamma = 30$ minutes, and $\Gamma = 2.5$ hours, then the maximum standard deviation of any $\delta \in \mathcal{D}$ exceeds the maximum empirical standard deviation by a factor of 2.5. Thus, \mathcal{D} contains extreme frequency deviation scenarios with unrealistically high variance.

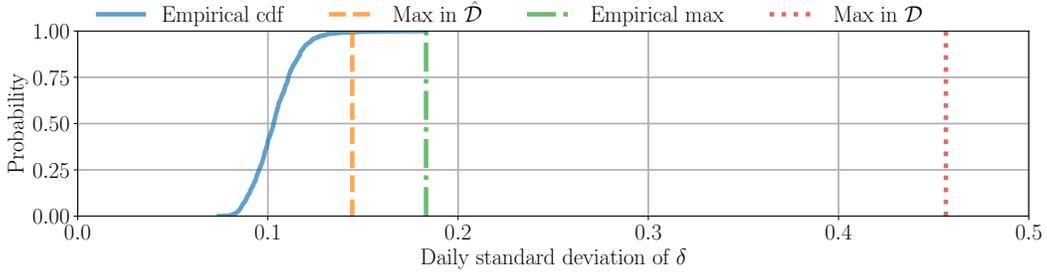


Figure 1. Empirical cumulative distribution function (cdf) of the daily standard deviation of δ and the maximum standard deviation of any scenario in $\hat{\mathcal{D}}$ or in \mathcal{D} .

The optimization model developed below not only involves the conservative uncertainty set \mathcal{D} compatible with the guidelines of the European Commission but also a smaller uncertainty set

$$\hat{\mathcal{D}} = \left\{ \delta \in \mathcal{L}(\mathcal{T}, [-1, 1]) : \int_{[t-\hat{\Gamma}]^+}^t |\delta(t')| dt' \leq \hat{\gamma} \forall t \in \mathcal{T} \right\}$$

parametrized by $\hat{\Gamma} \geq \Gamma$ and $\hat{\gamma} \leq \gamma$. This uncertainty set contains only frequency deviation scenarios that are likely to materialize under normal operating conditions. Note that $\hat{\mathcal{D}}$ is obtained from \mathcal{D} by inflating Γ to $\hat{\Gamma}$ and shrinking γ to $\hat{\gamma}$. By Remark 2, we may thus conclude that $\hat{\mathcal{D}}$ is indeed a subset of \mathcal{D} . While the pessimistic uncertainty set \mathcal{D} is used to enforce the stringent delivery guarantees imposed by the European Commission, the more optimistic uncertainty set $\hat{\mathcal{D}}$ is used to model a softer reachability guarantee for the terminal state-of-charge. In the numerical experiments we will set $\hat{\Gamma} = T = 1$ day and $\hat{\gamma} = \gamma = 30$ minutes. By Proposition 2, the variance of all frequency deviation scenarios in $\hat{\mathcal{D}}$ is therefore bounded above by $\Delta t/T = 1/48$. Empirically, this threshold exceeds the variance of the frequency deviation on 99.2% of all days in the years from 2017 to 2018.

Next, we discuss the uncertainty in the initial battery state-of-charge y_0 . Recall that y_0 is uncertain at the time when x^b and x^r are chosen because it depends on how much regulation energy must be provided until the beginning of the planning horizon. This quantity depends itself on uncertain frequency deviations that have not yet been revealed. We assume that the vehicle owner constructs two confidence intervals $\mathcal{Y}_0 = [y_0, \bar{y}_0]$ and $\hat{\mathcal{Y}}_0 = [\hat{y}_0, \hat{\bar{y}}_0]$ for y_0 , either taking into account all frequency deviations under which she must imperatively be able to deliver regulation power or only those frequency deviations that are likely to occur under normal operating conditions.

The only assumption we make about the uncertainty in the delivery price p^d is that the vehicle owner can reliably estimate the *expected* regulation price $p^r(t) = p^a(t) + \mathbb{E}[\tilde{\delta}(t)\tilde{p}^d(t)]$.

We are now ready to formalize the vehicle owner's decision problem for selecting the market decisions x^b and x^r . The primary objective is to minimize the expected cost

$$c(x^b, x^r) = \mathbb{E} \int_{\mathcal{T}} p^b(t) x^b(t) - \left(p^a(t) + \tilde{\delta}(t) \tilde{p}^d(t) \right) x^r(t) dt' = \int_{\mathcal{T}} p^b(t) x^b(t) - p^r(t) x^r(t) dt, \quad (3)$$

while ensuring that x^b and x^r are robustly feasible across all frequency deviation scenarios $\delta \in \mathcal{D}$ and initial battery states $y_0 \in \mathcal{Y}_0$. Mathematically, the charging rate $y^+(x^b(t), x^r(t), \delta(t))$, the discharging rate $y^-(x^b(t), x^r(t), \delta(t))$, and the battery state-of-charge $y(x^b, x^r, \delta, t, y_0)$ must therefore satisfy the robust constraints

$$\left. \begin{aligned} y^+(x^b(t), x^r(t), \delta(t)) &\leq \bar{y}^+(t), & y(x^b, x^r, \delta, y_0, t) &\leq \bar{y}, \\ y^-(x^b(t), x^r(t), \delta(t)) &\leq \bar{y}^-(t), & y(x^b, x^r, \delta, y_0, t) &\geq \underline{y} \end{aligned} \right\} \forall t \in \mathcal{T}, \forall \delta \in \mathcal{D}, \forall y_0 \in \mathcal{Y}_0.$$

As the vehicle owner continues to use the vehicle for driving and for offering grid services after the end of the planning horizon, the battery should end up in a state that is “*conducive to satisfactory future operations*” (Yeh, 1985). Consequently, the vehicle owner aims to steer $y(x^b, x^r, \delta, y_0, T)$ to a desirable state-of-charge y . We assume that the cost-to-go of any $y \in [\underline{y}, \bar{y}]$ is quantified by a convex and piecewise affine value function $\varphi(y) = \max_{n \in \mathcal{N}} \{a_n y + b_n\}$ determined by $a_n, b_n \in \mathbb{R}$ for all $n \in \mathcal{N} = \{1, \dots, N\}$. In Section 5, we will present a principled approach to calibrate these coefficients by solving a dynamic programming problem over multiple periods of length T . As y_0 and δ are uncertain, the terminal state-of-charge is also uncertain. To trade off present versus future costs, it is therefore reasonable to minimize $\varphi(y(x^b, x^r, \delta, y_0, T))$ in view of the worst of all scenarios $\delta \in \hat{\mathcal{D}}$ and $y_0 \in \hat{\mathcal{Y}}_0$. This can be achieved by adding the term $\max_{\delta \in \hat{\mathcal{D}}} \max_{y_0 \in \hat{\mathcal{Y}}_0} \varphi(y(x^b, x^r, \delta, y_0, T))$ to the objective function (3). In summary, the vehicle owner's decision problem can be cast as the following robust optimization problem with continuous (functional) uncertain parameters,

$$\begin{aligned} \min_{x^b, x^r \in \mathcal{X}} \quad & c(x^b, x^r) + \max_{\delta \in \hat{\mathcal{D}}, y_0 \in \hat{\mathcal{Y}}_0} \varphi(y(x^b, x^r, \delta, y_0, T)) \\ \text{s.t.} \quad & y^+(x^b(t), x^r(t), \delta(t)) \leq \bar{y}^+(t) \quad \forall \delta \in \mathcal{D}, \forall t \in \mathcal{T} \\ & y^-(x^b(t), x^r(t), \delta(t)) \leq \bar{y}^-(t) \quad \forall \delta \in \mathcal{D}, \forall t \in \mathcal{T} \\ & y(x^b, x^r, \delta, y_0, t) \leq \bar{y} \quad \forall \delta \in \mathcal{D}, \forall t \in \mathcal{T}, \forall y_0 \in \mathcal{Y}_0 \\ & y(x^b, x^r, \delta, y_0, t) \geq \underline{y} \quad \forall \delta \in \mathcal{D}, \forall t \in \mathcal{T}, \forall y_0 \in \mathcal{Y}_0, \end{aligned} \quad (\text{R})$$

where \mathcal{X} denotes the set of all functions in $\mathcal{L}(\mathcal{T}, \mathbb{R}_+)$ that are constant on the trading intervals. Using the conservative uncertainty sets \mathcal{D} and \mathcal{Y}_0 in the constraints ensures that the delivery guarantee

dictated by the European Commission can be fulfilled. Failing to fulfill this guarantee might lead to exclusion from the regulation market. In contrast, there are no drastic consequences of reaching an undesirable state-of-charge at time T . Hence, we use the less conservative uncertainty sets $\hat{\mathcal{D}}$ and $\hat{\mathcal{Y}}_0$ in the objective function, thereby ensuring that we end up in a desirable state-of-charge under all reasonably likely frequency deviation scenarios. The use of different uncertainty sets in the same model has previously been proposed in robust portfolio optimization (Zymler et al., 2011).

Recall from Proposition 1 that the function $y(x^b, x^r, \delta, y_0, t)$ is concave in the decision variables x^b and x^r . Upper bounds on this function thus constitute non-convex constraints. This implies that (R) represents a non-convex robust optimization problem with functional uncertain parameters. In general, such problems are severely intractable.

Remark 3 (Uncertain driving patterns). Although model (R) assumes deterministic driving patterns, it readily extends to uncertain driving times and distances. If it is only known that the vehicle will drive at some time within a prespecified interval, then the vehicle owner must not plan on exchanging any electricity with the grid during that interval. Similarly, if it is only known that the vehicle will drive some distance within a certain range, then the vehicle owner must plan with the low end of the range for the constraint on the maximum state-of-charge and with the high end of the range for the constraint on the minimum state-of-charge. The worst-case driving times and distances are thus independent of the vehicle owner’s decisions and can be determined *ex-ante*. \square

3 Time Discretization

In order to derive a lossless time discretization of the frequency deviation scenarios in problem (R), we assume from now on that the power demand for driving and the maximum charge and discharge power of the vehicle charger remain constant over the trading intervals. This assumption is justified because a vehicle that is both driving and parking in the same trading interval cannot offer constant market bids and is therefore unable to participate in the electricity market. Although the power demand for driving may fluctuate wildly, the battery state-of-charge cannot increase while the vehicle is driving, and therefore the power consumption for driving can be averaged over trading intervals without loss of generality. Note that we do *not* assume the frequency deviation scenarios δ to remain constant over the trading intervals. In practice δ may fluctuate on time scales of the order

of milliseconds, and averaging out the frequency deviations across a trading interval could result in a dangerous oversimplification of reality. This phenomenon is illustrated in the following example.

Example 1 (Risks of ignoring intra-period fluctuations). As the market decisions x^b and x^r , the power demand d and the charging limits \bar{y}^+ and \bar{y}^- are piecewise constant, one might be tempted to replace the frequency deviation signal δ with a piecewise constant signal obtained by averaging δ over the trading intervals. As we will see, however, averaging δ relaxes the battery state-of-charge constraints. Decisions x^b and x^r that are *infeasible* under the true signal may therefore appear to be feasible under the averaged signal. Hence, replacing the true signal with the averaged signal could make it impossible for the vehicle owner to honor her market commitments. As a simple example, assume that $x^b(t) = 0$ and $x^r(t) = x_1^r > 0$ are constant and that the true frequency deviation signal averages to 0 over the first trading interval $[0, \Delta t]$, that is, $\frac{1}{\Delta t} \int_0^{\Delta t} \delta(t) dt = 0$. The left chart of Figure 2 visualizes two such signals, which display a small and a high total variation and are denoted by $\delta^{(1)}$ and $\delta^{(2)}$, respectively. The constant signal equal to their (vanishing) average over $[0, \Delta t]$ is denoted by $\delta^{(3)}$. If $\delta^{(1)}$ reflects reality but is incorrectly replaced with $\delta^{(3)}$, we are led to believe that the state-of-charge will remain constant at y_0 . In reality, however, the battery dissipates the amount $\Delta\eta x_1^r \Delta t / 2$ of energy over the first trading interval, where $\Delta\eta = \frac{1}{\eta^-} - \eta^+ \geq 0$, and the state-of-charge temporarily rises above y_0 by $\eta^+ x_1^r \Delta t / 2$. If $\delta^{(2)}$ reflects reality, on the other hand, then the repeated charging and discharging of the battery still dissipates energy. See the right chart of Figure 2 for a visualization. While scenario $\delta^{(1)}$ is contrived for maximum impact, scenario $\delta^{(2)}$ rapidly fluctuates around 0 and thus captures a stylized fact that one would expect to see in reality. This example suggests that finding the minimum or the maximum of the state-of-charge over the entire planning horizon and over all signals $\delta \in \mathcal{D}$ should be non-trivial because intra-period fluctuations *do* matter. As a further complication, note that the constraints of the uncertainty set \mathcal{D} couple the frequency deviations across time. \square

We will now argue that, in spite of Example 1, \mathcal{D} and $\hat{\mathcal{D}}$ can be restricted to contain only piecewise constant frequency deviation signals *without* relaxing problem (R). To formalize the reasoning about piecewise constant functions, we introduce a lifting operator $L : \mathbb{R}^K \rightarrow \mathcal{L}(\mathcal{T}, \mathbb{R})$ that maps any vector $\mathbf{v} \in \mathbb{R}^K$ to a piecewise constant function $L\mathbf{v}$ with K pieces defined through $(L\mathbf{v})(t) = v_k$ if $t \in \mathcal{T}_k$, $k \in \mathcal{K} = \{1, \dots, K\}$. We also introduce the adjoint operator $L^\dagger : \mathcal{L}(\mathcal{T}, \mathbb{R}) \rightarrow \mathbb{R}^K$ that maps any function $w \in \mathcal{L}(\mathcal{T}, \mathbb{R})$ to a K -dimensional vector $L^\dagger w$ defined through $(L^\dagger w)_k = \frac{1}{\Delta t} \int_{\mathcal{T}_k} w(t) d(t)$ for

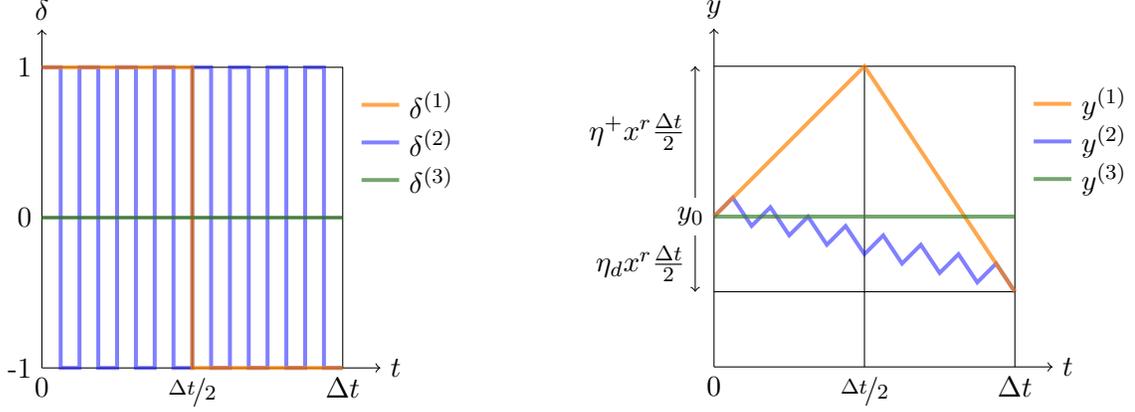


Figure 2. Frequency deviation signals (left) and their state-of-charge trajectories (right).

all $k \in \mathcal{K}$. Note that L and L^\dagger are indeed adjoint to each other because $\int_{\mathcal{T}} (L\mathbf{v})(t)w(t) dt = \mathbf{v}^\top L^\dagger(w)$ for all $\mathbf{v} \in \mathbb{R}^K$ and $w \in \mathcal{L}(\mathcal{T}, \mathbb{R})$. Mathematically, we impose from now on the following assumption.

Assumption 1. The functions d , \bar{y}^+ and \bar{y}^- are piecewise constant, that is, there exist $\mathbf{d}, \bar{\mathbf{y}}^+, \bar{\mathbf{y}}^- \in \mathbb{R}^K$ such that $d = L\mathbf{d}$, $\bar{y}^+ = L\bar{\mathbf{y}}^+$ and $\bar{y}^- = L\bar{\mathbf{y}}^-$.

Next, we introduce a discretized uncertainty set

$$\mathcal{D}_{\mathcal{K}} = \left\{ \boldsymbol{\delta} \in [-1, 1]^K : \sum_{l=1+[k-\Gamma/\Delta t]^+}^k |\delta_l| \leq \frac{\gamma}{\Delta t} \forall k \in \mathcal{K} \right\}$$

reminiscent of \mathcal{D} , where $\Gamma/\Delta t$ and $\gamma/\Delta t$ count the trading intervals within a regulation cycle and an activation period, respectively. Similarly, we define a smaller discretized uncertainty set $\hat{\mathcal{D}}_{\mathcal{K}} \subseteq \mathbb{R}^K$ reminiscent of $\hat{\mathcal{D}}$, which is obtained from $\mathcal{D}_{\mathcal{K}}$ by replacing Γ with $\hat{\Gamma}$ and γ with $\hat{\gamma}$. In the remainder we impose the following divisibility assumption.

Assumption 2. The parameters Γ , γ , $\hat{\Gamma}$ and $\hat{\gamma}$ are (positive) multiples of Δt .

The discretized uncertainty sets are of interest because of the following proposition.

Proposition 3. *The following statements hold.*

(i) $LD_{\mathcal{K}} \subseteq \mathcal{D}$ and $L^\dagger \mathcal{D} = \mathcal{D}_{\mathcal{K}}$.

(ii) $LD_{\mathcal{K}}^+ \subseteq \mathcal{D}^+$ and $L^\dagger \mathcal{D}^+ = \mathcal{D}_{\mathcal{K}}^+$.

Next, we define the finite-dimensional feasible set $\mathcal{X}_{\mathcal{K}} = L^\dagger \mathcal{X}$. As \mathcal{X} contains only piecewise constant functions, we have $\mathcal{X} = L\mathcal{X}_{\mathcal{K}}$. We further define the cost function $c_{\mathcal{K}}(\mathbf{x}^b, \mathbf{x}^r) = c(L\mathbf{x}^b, L\mathbf{x}^r)$, which is linear in $\mathbf{x}^b \in \mathbb{R}^K$ and $\mathbf{x}^r \in \mathbb{R}^K$. In addition, for any $k \in \mathcal{K}$ we define the function

$$\begin{aligned} y_k(\mathbf{x}^b, \mathbf{x}^r, \boldsymbol{\delta}, y_0) &= y(L\mathbf{x}^b, L\mathbf{x}^r, L\boldsymbol{\delta}, y_0, k\Delta t) \\ &= y_0 + \Delta t \sum_{l=1}^k \eta^+ y^+(x_l^b, x_l^r, \delta_l) - \frac{1}{\eta^-} y^-(x_l^b, x_l^r, \delta_l) - d_l, \end{aligned} \quad (4)$$

which represents the battery state-of-charge at the end of period k under the assumption that both the market bids and the frequency deviations are piecewise constant.

Proposition 4. *Holding all other factors fixed, $y_k(\mathbf{x}^b, \mathbf{x}^r, \boldsymbol{\delta}, y_0)$ is concave nondecreasing in \mathbf{x}^b , concave in \mathbf{x}^r , concave nondecreasing in $\boldsymbol{\delta}$, and linear nondecreasing in y_0 for any $k \in \mathcal{K}$.*

We are now ready to define the discrete-time counterpart of the robust optimization problem (R).

$$\begin{aligned} \min_{\mathbf{x}^b, \mathbf{x}^r \in \mathcal{X}_{\mathcal{K}}} \quad & c_{\mathcal{K}}(\mathbf{x}^b, \mathbf{x}^r) + \max_{\boldsymbol{\delta} \in \hat{\mathcal{D}}_{\mathcal{K}}, y_0 \in \hat{\mathcal{Y}}_0} \varphi(y_{\mathcal{K}}(\mathbf{x}^b, \mathbf{x}^r, \boldsymbol{\delta}, y_0)) \\ \text{s.t.} \quad & y^+(x_k^b, x_k^r, \delta_k) \leq \bar{y}_k^+ \quad \forall \boldsymbol{\delta} \in \mathcal{D}_{\mathcal{K}}, \forall k \in \mathcal{K} \\ & y^-(x_k^b, x_k^r, \delta_k) \leq \bar{y}_k^- \quad \forall \boldsymbol{\delta} \in \mathcal{D}_{\mathcal{K}}, \forall k \in \mathcal{K} \\ & y_k(\mathbf{x}^b, \mathbf{x}^r, \boldsymbol{\delta}, y_0) \leq \bar{y} \quad \forall \boldsymbol{\delta} \in \mathcal{D}_{\mathcal{K}}, \forall k \in \mathcal{K} \cup \{0\}, \forall y_0 \in \mathcal{Y}_0 \\ & y_k(\mathbf{x}^b, \mathbf{x}^r, \boldsymbol{\delta}, y_0) \geq \underline{y} \quad \forall \boldsymbol{\delta} \in \mathcal{D}_{\mathcal{K}}, \forall k \in \mathcal{K} \cup \{0\}, \forall y_0 \in \mathcal{Y}_0 \end{aligned} \quad (\text{R}_{\mathcal{K}})$$

Unlike the original problem (R), the discrete-time counterpart (R_κ) constitutes a standard robust optimization problem that involves only finite-dimensional uncertain parameters. For this reason, there is hope that (R_κ) is easier to solve than (R).

Theorem 1 (Lossless time discretization). *The problems (R) and (R_κ) are equivalent.*

The equivalence of (R) and (R_κ) is perhaps surprising in view of Example 1. It means that the worst-case frequency deviation scenarios are piecewise constant even though intra-period fluctuations matter. Theorem 1 is proved by showing that the four robust constraints in (R) with functional uncertainties are equivalent to the corresponding robust constraints in (R_κ) with vectorial uncertainties and that the worst-case terminal cost functions in (R) and (R_κ) coincide. For example, the equivalence of the first (second) robust constraints in (R) and (R_κ) follows from the observation that, for any fixed $t \in \mathcal{T}$, the left hand side of the first (second) constraint in (R) is

maximized by a scenario $\delta \in \mathcal{D}$ with $\delta(t) = 1$ ($\delta(t) = -1$), which exists by the definition of \mathcal{D} . The last two robust constraints in (R) are nonlocal as they depend on the entire frequency deviation scenario δ and not only on its value at a particular time. Thus, they are significantly more intricate. The robust upper bound on the state-of-charge can be reformulated as an upper bound on $\max_{t \in \mathcal{T}} \max_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, t)$. By using Propositions 1 and 3, one can then show that the maximum over $t \in \mathcal{T}$ must be attained at $t = k\Delta t$ for some $k \in \mathcal{K} \cup \{0\}$ and that for any such t the state-of-charge can be expressed as an integral of δ against a piecewise constant function. Thus, averaging δ across the trading intervals has no impact on the state-of-charge, which in turn allows us to focus on piecewise constant scenarios without restricting generality. The robust lower bound on the state-of-charge in (R) can be reformulated as a lower bound on $\min_{t \in \mathcal{T}} \min_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, t)$, which appears to be intractable because the optimization problem over δ minimizes a concave function over a convex feasible set and is therefore non-convex. Classical robust optimization provides no general recipe for handling such constraints even if the uncertain parameters are finite-dimensional, and state-of-the-art research settles for deriving approximations (Roos et al., 2020). By exploiting a continuous total unimodularity property of the uncertainty set \mathcal{D} facilitated by Assumption 2, we first prove that the minimum of $y(x^b, x^r, \delta, y_0, t)$ over \mathcal{D} is attained by a frequency deviation trajectory that takes only values in $\{-1, 0\}$. Next, we demonstrate that there exists an affine function of δ that matches $y(x^b, x^r, \delta, y_0, t)$ for all trajectories $\delta \in \mathcal{D}$ valued in $\{-1, 0\}$ and for all $t \in \mathcal{T}$. In the language of robust optimization, the state-of-charge $y(x^b, x^r, \delta, y_0, t)$ can be viewed as an analysis variable that adapts to the uncertainty δ , and the corresponding affine function constitutes a decision rule approximating $y(x^b, x^r, \delta, y_0, t)$. Decision rule approximations almost invariably introduce approximation errors (Ben-Tal et al., 2009, § 14). However, the affine decision rule proposed here is error-free because it coincides with $y(x^b, x^r, \delta, y_0, t)$ for all scenarios $\delta \in \mathcal{D}$ valued in $\{-1, 0\}$ that may attain the worst case in $\min_{t \in \mathcal{T}} \min_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, t)$. Using this decision rule in an elaborate sensitivity analysis, we can finally prove that the minimum over $t \in \mathcal{T}$ must be attained at $t = k\Delta t$ for some $k \in \mathcal{K} \cup \{0\}$ and that for any such t the state-of-charge can be expressed as an integral of δ against a piecewise constant function. Thus, we may focus again on piecewise constant scenarios without restricting generality. The full proof of Theorem 1 can be found in Appendix C.

The new robust optimization techniques developed to prove Theorem 1 are of independent interest as they provide exact tractable reformulations for certain adjustable robust optimization

problems with functional or vectorial uncertain parameters, where the embedded optimization problems over the uncertainty realizations are non-convex. We also note that the embedded optimization problems over $\delta \in \mathcal{D}$ in problem (R) can be viewed as variants of the so-called *separated continuous linear programs* introduced by Anderson et al. (1983). The proof of Theorem 1 shows that these problems are solved by piecewise constant frequency deviation scenarios that can be computed efficiently, thereby extending the purely existential results by Pullan (1995). Our results are also orthogonal to those by Ghate (2020), who proves that *robust* separated continuous linear programs with budget uncertainty sets are equivalent to standard separated continuous linear programs.

Even though the non-convex robust optimization problem (R) with functional uncertainty admits a lossless time discretization, its discrete-time counterpart $(R_{\mathcal{K}})$ still constitutes a non-convex robust optimization problem and thus appears to be hard. In the next section, however, we will show that (R) can be reformulated as a tractable linear program by exploiting its structural properties.

4 Linear Programming Reformulation

In order to establish the tractability of the non-convex robust optimization problem (R), it is useful to reformulate its time discretization $(R_{\mathcal{K}})$ as the following *linear* robust optimization problem, where all constraint functions are bilinear in the decision variables and the uncertain parameters.

$$\begin{aligned}
\min \quad & c_{\mathcal{K}}(\mathbf{x}^b, \mathbf{x}^r) + z \\
\text{s.t.} \quad & \mathbf{x}^b, \mathbf{x}^r \in \mathcal{X}_{\mathcal{K}}, \mathbf{m} \in \mathbb{R}^K \\
& x_k^r + x_k^b \leq \bar{y}_k^+, \quad x_k^r - x_k^b \leq \bar{y}_k^- \quad \forall k \in \mathcal{K} \\
& m_k - \eta^+ x_k^r \geq 0, \quad m_k - \frac{1}{\eta^-} x_k^r + \Delta \eta x_k^b \geq 0 \quad \forall k \in \mathcal{K} \\
& \bar{y}_0 + \Delta t \sum_{l=1}^k \eta^+ (x_l^b + \delta_l x_l^r) - d_l \leq \bar{y} \quad \forall \delta \in \mathcal{D}_{\mathcal{K}}^+, \forall k \in \mathcal{K} \cup \{0\} \\
& \underline{y}_0 + \Delta t \sum_{l=1}^k \eta^+ x_l^b - m_l \delta_l - d_l \geq \underline{y} \quad \forall \delta \in \mathcal{D}_{\mathcal{K}}^+, \forall k \in \mathcal{K} \cup \{0\} \\
& \hat{y}_0 + \Delta t \sum_{k=1}^K \eta^+ (x_k^b + \delta_k x_k^r) - d_k \leq \frac{z - b_n}{a_n} \quad \forall \delta \in \hat{\mathcal{D}}_{\mathcal{K}}^+, \forall n \in \mathcal{N}_+ \\
& \hat{y}_0 + \Delta t \sum_{k=1}^K \eta^+ x_k^b - m_k \delta_k - d_k \geq \frac{z - b_n}{a_n} \quad \forall \delta \in \hat{\mathcal{D}}_{\mathcal{K}}^+, \forall n \in \mathcal{N}_- \\
& z \geq b_n \quad \forall n \in \mathcal{N}_0
\end{aligned} \tag{R'_{\mathcal{K}}}$$

Here, $\Delta\eta = \frac{1}{\eta} - \eta^+ \geq 0$ is used as a shorthand for the reduction in the battery state-of-charge resulting from first charging and then discharging one unit of energy as seen from the grid. In addition, we set $\mathcal{N}_+ = \{n \in \mathcal{N} : a_n > 0\}$, $\mathcal{N}_- = \{n \in \mathcal{N} : a_n < 0\}$ and $\mathcal{N}_0 = \mathcal{N} \setminus (\mathcal{N}_+ \cup \mathcal{N}_-)$.

Theorem 2 (Lossless linearization). *The problems $(R_{\mathcal{K}})$ and $(R'_{\mathcal{K}})$ are equivalent.*

The proof of Theorem 2 critically relies on the exact affine decision rule approximation discovered in the proof of Theorem 1. Note that the linear robust optimization problem $(R'_{\mathcal{K}})$ still appears to be difficult because each robust constraint must hold for all frequency deviation scenarios in an uncountable uncertainty set $\mathcal{D}_{\mathcal{K}}^+$ or $\hat{\mathcal{D}}_{\mathcal{K}}^+$ and therefore corresponds to a continuum of ordinary linear constraints. Fortunately, standard robust optimization theory (Ben-Tal et al., 2004; Bertsimas and Sim, 2004) allows us to reformulate $(R'_{\mathcal{K}})$ as the tractable linear program

$$\begin{aligned}
& \min && c_{\mathcal{K}}(\mathbf{x}^b, \mathbf{x}^r) + z \\
& \text{s.t.} && \mathbf{x}^b, \mathbf{x}^r \in \mathcal{X}_{\mathcal{K}}, \quad z \in \mathbb{R}, \quad \mathbf{m}, \boldsymbol{\lambda}^+, \boldsymbol{\lambda}^-, \boldsymbol{\theta}^+, \boldsymbol{\theta}^- \in \mathbb{R}_+^K, \quad \boldsymbol{\Lambda}^+, \boldsymbol{\Lambda}^-, \boldsymbol{\Theta}^+, \boldsymbol{\Theta}^- \in \mathbb{R}_+^{K \times K} \\
& && x_k^r + x_k^b \leq \bar{y}_k^+, \quad x_k^r - x_k^b \leq \bar{y}_k^- \quad \forall k \in \mathcal{K} \\
& && m_k \geq \eta^+ x_k^r, \quad m_k \geq \frac{1}{\eta^-} x_k^r - \Delta\eta x_k^b \quad \forall k \in \mathcal{K} \\
& && \sum_{l=1}^k \Delta t \left(\eta^+ x_l^b + \Lambda_{k,l}^+ - d_l \right) + \gamma \Theta_{k,l}^+ \leq \bar{y} - \bar{y}_0 \quad \forall k \in \mathcal{K} \cup \{0\} \\
& && \sum_{l=1}^k \Delta t \left(\eta^+ x_l^b - \Lambda_{k,l}^- - d_l \right) - \gamma \Theta_{k,l}^- \geq \underline{y} - \underline{y}_0 \quad \forall k \in \mathcal{K} \cup \{0\} \\
& && \sum_{k \in \mathcal{K}} \Delta t \left(\eta^+ x_k^b + \lambda_k^+ - d_k \right) + \hat{\gamma} \theta_k^+ \leq \frac{z - b_n}{a_n} - \hat{y}_0 \quad \forall n \in \mathcal{N}_+ \\
& && \sum_{k \in \mathcal{K}} \Delta t \left(\eta^+ x_k^b - \lambda_k^- - d_k \right) - \hat{\gamma} \theta_k^- \geq \frac{z - b_n}{a_n} - \hat{y}_0 \quad \forall n \in \mathcal{N}_- \\
& && b_n \leq z \quad \forall n \in \mathcal{N}_0 \\
& && \Lambda_{k,l}^+ + \sum_{i=l}^{I(k,l)} \Theta_{k,i}^+ \geq \eta^+ x_l^r, \quad \Lambda_{k,l}^- + \sum_{i=l}^{I(k,l)} \Theta_{k,i}^- \geq m_l \quad \forall k, l \in \mathcal{K} : l \leq k \\
& && \lambda_k^+ + \sum_{i=k}^{\hat{I}(K,k)} \theta_i^+ \geq \eta^+ x_k^r, \quad \lambda_k^- + \sum_{i=k}^{\hat{I}(K,k)} \theta_i^- \geq m_k \quad \forall k \in \mathcal{K},
\end{aligned} \tag{R''_{\mathcal{K}}}$$

where $I(k, l) = \min\{k, l + \Gamma/\Delta t - 1\}$ and $\hat{I}(k, l) = \min\{k, l + \hat{\Gamma}/\Delta t - 1\}$.

Theorem 3 (Linear programming reformulation). *The problems $(R'_{\mathcal{K}})$ and $(R''_{\mathcal{K}})$ are equivalent.*

The conversion of the robust optimization problem $(R'_{\mathcal{K}})$ to the linear program $(R''_{\mathcal{K}})$ comes at the expense of introducing $4K^2 + 4K$ dual variables. For a daily planning horizon with half-hourly resolution, this amounts to introducing 9,408 additional continuous variables. Overall, the linear program $(R''_{\mathcal{K}})$ involves $4K^2 + 7K + 1$ variables and $K^2 + 9K + N + 2$ constraints, that is, its size scales quadratically with K . In conjunction, Theorems 1–3 imply that the non-convex robust optimization problem (R) with continuous uncertain parameters can be reduced without any loss to the tractable linear program $(R''_{\mathcal{K}})$, which is amenable to efficient numerical solution with state-of-the-art linear programming solvers such as CPLEX or Gurobi.

Remark 4 (Robustification reduces complexity). A striking property of the robust optimization model (R) is that it is much *easier* to solve than the underlying deterministic model, which would assume precise knowledge of the frequency deviation scenario δ . Indeed, the textbook formulation of the deterministic model requires continuous decision variables to represent $y^+(x^b(t), x^r(t), \delta(t))$ and $y^-(x^b(t), x^r(t), \delta(t))$ and a binary decision variable to model their complementarity for every $t \in \mathcal{T}$ (Taylor, 2015, p. 85). This results in a large-scale mixed-integer linear program even if \mathcal{T} is discretized. In contrast, the robust optimization model (R) is equivalent to the tractable linear program $(R''_{\mathcal{K}})$. To our best knowledge, we have thus discovered the first practically interesting class of optimization problems that become dramatically easier through robustification. \square

5 Multi-Stage Extensions

Model (R) described in Section 2 looks only one day ahead and accounts for the future usage of the vehicle only through the cost-to-go function φ . We will now show that this static decision problem readily extends to a dynamic model that looks H days into the future. For ease of exposition, we assume that the market bids for day $h \in \mathcal{H} = \{0, \dots, H - 1\}$ are due at midnight of the previous day, which implies that the initial state-of-charge y_h on day h is precisely known. However, this assumption can be relaxed. A vehicle owner aiming to minimize worst-case expected costs (where the expectation is taken with respect to the prices, and the worst-case is taken with respect to the frequency deviation scenarios in $\hat{\mathcal{D}}$) across H days solves the robust dynamic program

$$\varphi_h(y_h) = \min_{(x_h^b, x_h^r) \in \mathcal{X}_{\mathcal{R}}(y_h)} c(x_h^b, x_h^r) + \max_{\delta_h \in \hat{\mathcal{D}}} \varphi_{h+1}(y(x_h^b, x_h^r, \delta_h, y_h, T)) \quad \forall h \in \mathcal{H}. \quad (5)$$

Here, $\mathcal{X}_R(y_h)$ denotes the feasible set of the single-stage problem (R) with uncertainty sets \mathcal{D} and $\mathcal{Y}_0 = \{y_h\}$. In principle, $\mathcal{X}_R(y_h)$ could be empty for some y_h . Under the reasonable assumption that each night the vehicle is plugged in long enough for the battery to be fully charged, however, $\mathcal{X}_R(y_h)$ is guaranteed to be non-empty. In this case, problem (5) remains feasible on all days $h \in \mathcal{H}$ irrespective of previous market decisions and even when facing a frequency deviation scenario $\delta \notin \hat{\mathcal{D}}$. The cost-to-go functions φ_h , $h \in \mathcal{H}$, are defined recursively through (5). In practice, it proves useful to initialize the recursion by $\varphi_H(y_H) = p^*|y_H - y^*|$ for some target state-of-charge y^* and penalty parameter $p^* \geq 0$, which can be calibrated to historical data via cross-validation (see Section 6.2).

Proposition 5. *The cost-to-go function φ_h is convex and piecewise affine for every $h \in \mathcal{H}$.*

By Proposition 5, problem (5) is structurally equivalent to problem (R) for any $h \in \mathcal{H}$. Theorems 1, 2 and 3 thus imply that (5) can be recast as a tractable linear program whose right hand side coefficients depend affinely on y_h . Computing the function value and an arbitrary subgradient of φ_h at a fixed y_h is therefore tantamount to solving a linear program. This insight suggests that convex piecewise affine bounds on the cost-to-go function φ_h can be constructed by solving a series of linear programs. If φ_{h+1} is precisely known for some $h \in \mathcal{H}$, for example, then a coarse upper bound is obtained by evaluating φ_h at \underline{y} and \bar{y} and by linearly interpolating the two function values. A coarse lower bound is given by the pointwise maximum of the two tangents of φ_h constructed from the function values and subgradients at \underline{y} and \bar{y} . The difference between these coarse bounds is maximal at the kink of the lower bound. To improve both bounds, we can then break the interval $[\underline{y}, \bar{y}]$ apart at the kink and construct separate upper and lower bounds on the two resulting subintervals by repeating the above procedure. Iteratively partitioning the subinterval on which the gap between the bounds is maximal yields increasingly tight convex piecewise affine bounds that approximate φ_h uniformly on $[\underline{y}, \bar{y}]$ to any precision. In each iteration one has to solve a linear program akin to $(R_{\mathcal{K}}'')$ in order to compute the value and a subgradient of φ_h at a new anchor point. Note that the outlined procedure remains applicable if φ_{h+1} in (5) is replaced with a convex piecewise affine bound, in which case one has only access to an *inexact* oracle for the values and subgradients of φ_h . In this case the approximation errors accumulate over the dynamic programming recursions.

In reality, the market bids for day h are due around noon and not at midnight of day $h - 1$. As a consequence, the exact state-of-charge y_h at the beginning of day h is unknown at stage h , that is, at the time when (x_h^b, x_h^r) must be chosen. Instead, only the confidence intervals $[\hat{y}_0, \hat{\hat{y}}_0]$

and $[y_0, \bar{y}_0]$ for y_h are available, which can be constructed from the current state-of-charge and from (x_{h-1}^b, x_{h-1}^r) ; see Section 2. In a more realistic multi-stage model, the information available at stage h is therefore encoded by the state vector $(\hat{y}_0, \hat{\bar{y}}_0, \underline{y}_0, \bar{y}_0)$, and thus the cost-to-go functions have four arguments, while the state transition functions have four components. By generalizing Theorems 1, 2 and 3 as well as Proposition 5, one can still show that all cost-to-go functions are convex and piecewise affine and that their values and subgradients can be computed by solving linear programs akin to $(R_{\mathcal{K}}'')$. While tedious, the proofs of these generalized results require no fundamentally new ideas and are thus omitted for brevity. As the cost-to-go functions have four arguments, the construction of convex piecewise affine upper and lower bounds is more challenging but still possible via the robust dual dynamic programming algorithm by Georghiou et al. (2019).

Numerical experiments in Section 6.3 suggest that the benefits of solving a multi-stage instead of a single-stage problem are negligible in practice because electric vehicles can always be charged overnight. We emphasize, however, that the robust optimization models and techniques developed in this paper may also be useful to optimize the operation of other energy storage devices that are characterized by slower dynamics and therefore necessitate a proper multi-stage approach.

6 Numerical Experiments

In the following, we first describe how the vehicle owner’s decision problem is parametrized from data, and we explain the backtesting procedure that is used to assess the performance of a given bidding strategy. Next, we present numerical results and discuss policy implications. All experiments are run on an Intel i7-6700 CPU with 3.40GHz clock speed and 64GB of RAM. All linear programs are solved with GUROBI 9.1.2 using its PYTHON interface. In order to ensure the reproducibility of our experiments, we provide links to all data sources and make our code available at www.github.com/lauinger/reliable-frequency-regulation-through-vehicle-to-grid.

6.1 Model Parametrization

The French transmission system operator (RTE) publishes availability and delivery prices and frequency measurements.⁴ There have been two policy changes in frequency regulation since 2015. While the availability prices were historically kept constant throughout the year, they change on

⁴<http://clients.rte-france.com/>

a weekly basis since mid-January 2017 and on a daily basis since July 2019. At this point, the pricing mechanism also changed from a pay-as-bid auction to a clearing price auction. The average availability price over all years from 2015 to 2019 amounts to 0.8cts/kWh, but the yearly average *decreased* in 2017 and 2018, and *increased* again in 2019 to pre-2017 levels. For all practical purposes we may assume that the expected regulation price $p^r(t) = p^a(t) + \mathbb{E}[\tilde{\delta}(t)\tilde{p}^d(t)]$ coincides with the availability price $p^a(t)$ because the realized regulation price $p^a(t) + \delta(t)p^d(t)$ oscillates rapidly around the availability price $p^a(t)$ due to intra-day fluctuations of the frequency-adjusted delivery prices; see Figure 7 in Appendix A. In fact, $\delta(t)p^d(t)$ empirically averages to $-2.36 \cdot 10^{-5}\text{€}$ over all 10s intervals from from 2015 to 2019. We further identify the utility prices $p^b(t)$ with the residential electricity prices charged by Electricité de France (EDF), the largest European electricity provider. These prices exhibit six different levels corresponding to peak- and off-peak hours on high, medium and low price days. High price days can occur exclusively on work days between November and March, whereas medium price days can occur on all days except Sundays. Low-price days can occur year-round. The peak hours are defined as the hours from 6 am to 10 pm on work days, and all the other hours are designated as off-peak hours. The prices corresponding to each type of day and hour are regulated and published in the official French government bulletin.⁵ Over the past five years, these prices have not changed more than three times per year. On each day, RTE announces the next day’s price levels by 10:30 am. The average utility price over the years from 2015 to 2019 amounts to 14cts/kWh and thus exceeds the average availability price by an order of magnitude.

When simulating the impact of the market decisions on the battery state-of-charge, it is important to track the frequency signal with a high time resolution. In fact, the European Commission (2017) requires regulation providers to adjust the power flow between the battery and the grid every ten seconds in order to ensure that it closely matches $x^b(t) + \delta(t)x^r(t)$ for all $t \in \mathcal{T}$. This means that regulation providers need to measure the frequency deviation $\delta(t)$ at least every ten seconds. Hence, we use a sampling rate of 100mHz when simulating the impact of the market decisions on the battery state-of-charge. This contrasts with previous studies, which used sampling rates below 20mHz (Han et al., 2010; Sortomme and El-Sharkawi, 2012; Donadee and Ilić, 2014; Wu and Sioshansi, 2019). Recall from Section 2 that the frequency deviation $\delta(t)$ depends on the nominal grid frequency $f_0 = 50\text{Hz}$ and on the normalization constant $\Delta f = 200\text{mHz}$. Moreover, recall that

⁵Journal Officiel de la République Française: <https://www.legifrance.gouv.fr/>

the uncertainty set \mathcal{D} is parametrized by $\gamma = 0.5\text{h}$ and $\Gamma = 2.5\text{h}$ in order to respect the delivery guarantee rules prescribed by the European Commission, while the less conservative uncertainty set $\hat{\mathcal{D}}$ is parametrized by $\hat{\gamma} = 0.5\text{h}$ and $\hat{\Gamma} = 24\text{h}$ as described in the discussion of Figure 1.

The vehicle data is summarized in Table 2 in Appendix A. The chosen parameter values are representative for commercially available midrange vehicle-to-grid-capable electric vehicles such as the 2018 Nissan Leaf, and they are in line with experimental measurements of charging and discharging efficiencies. For example, Apostolaki-Iosifidou et al. (2017) find that charging and discharging efficiencies η^+ and η^- , respectively, vary between 64% and 88% for a LiPF₆ cobalt battery with a nominal voltage of 345V and a nominal capacity of 106Ah. We assume that the vehicle owner reserves the time windows from 7 am to 9 am and from 5 pm to 7 pm on work days and from 8 am to 8 pm on weekend days and public holidays for driving. At all other times, the car is connected to a bidirectional charging station. We also assume that the car’s yearly mileage amounts to 10,000km, which approximately matches the French average of 13,000km (Commissariat général au développement durable, 2010). Hence, the car travels about 27km per day on average. Even though this distance is easily covered within one hour, it makes sense to reserve extended time slots for driving. Indeed, the vehicle owner may not be able to (nor wish to) pinpoint the exact driving times one day in advance. It may also be impossible to find bidirectional charging stations on weekend trips. With a standard vehicle efficiency of 0.2kWh/km, the car thus consumes 2,000kWh per year.

6.2 Backtesting Procedure and Baseline Strategy

In our experiments, we assess the performance of different bidding strategies over different test datasets covering one of the years between 2015 and 2019. A bidding strategy is any procedure that computes on each day at noon a pair of market decisions x^b and x^r for the following day. We call a strategy *non-anticipative* if it determines the market decisions using only information observed in the past. In addition, we call a strategy *feasible* if it allows the vehicle owner to honor all market commitments for all frequency deviation scenarios within the uncertainty set \mathcal{D} .

To measure the profit generated by a particular strategy over one year of test data, we use the following backtesting procedure. On each day at noon we compute the market decisions for the following day. We then use the actual frequency deviation data between noon and midnight and the market decisions for the current day to calculate the true battery state-of-charge at midnight.

Next, we use the frequency deviation data of the following day to calculate the revenue from selling regulation power to the TSO, which is subtracted from the cost of buying electricity for charging the battery. If the strategy is infeasible and the vehicle owner is not able to deliver all promised regulation power even though the realized frequency deviation trajectory falls within the uncertainty set \mathcal{D} , then she pays a penalty. The penalty at time t is set to $k_{\text{pen}} \cdot p^a(t) \cdot (x^r(t) - x_d^r(t))$, where $x_d^r(t)$ denotes the maximum amount of regulation power that could have been offered without risking an infeasibility, and the penalty factor k_{pen} ranges from 3 to 10.⁶ Repeated offenses may even lead to market exclusion. For simplicity, in each experiment we either assume that the vehicle owner pays a penalty corresponding to a fixed value of k_{pen} for every offense or is excluded from the regulation market directly upon the first offense. If the battery is depleted during a trip, any missing energy needed for driving is acquired at a high price p^y from a public fast charging station. We assume that p^y accounts for the price of energy as well as for the opportunity cost of the time lost in driving to the charging station and waiting to be serviced. In our experiments we set p^y either to 0.75€/kWh (which corresponds to typical energy prices offered by the European fast charging network Ionity⁷), to 7.5€/kWh or to 75€/kWh. The procedure described above is repeated on each day, and the resulting daily profits are accumulated over the entire test dataset.

Our baseline strategy is to determine the next day’s market decisions by solving the robust optimization problem (R) with terminal cost function $\varphi(y) = p^*|y - y^*|$, where the calibration of y^* and p^* is described below. Thus, (R) is equivalent to an instance of the linear program (R $''_{\mathcal{K}}$). This problem is updated on each day because the driving pattern d as well as the market prices p^b and p^r change, and because the uncertainty sets \mathcal{Y}_0 and $\hat{\mathcal{Y}}_0$ for the state-of-charge at midnight depend on the state-of-charge at noon and on the market commitments between noon and midnight that were chosen one day earlier. The baseline strategy is feasible thanks to the robust constraints in (R), which ensure that regulation power can be provided for all frequency deviation scenarios in \mathcal{D} .

The parameters p^* and y^* are kept constant throughout each backtest. Specifically, we set $p^* = \frac{3+k}{40}$ €/kWh for some $k = 1, \dots, 9$ and $y^* = (\frac{\bar{y}+y}{2} + l)$ kWh for some $l = 0, \dots, 5$. Every tuple (p^*, y^*) encodes a different bidding strategy. Given a training dataset comprising one year of frequency measurements and market prices, we compute the cumulative profit of each strategy via the backtesting procedure outlined above, and we choose the tuple (p^*, y^*) that corresponds to the

⁶For example, Réseau de transport d’électricité (2017a) sets $k_{\text{pen}} = 5$.

⁷<https://ionity.eu/en/>

Table 1. Calibration of p^* and y^* .

Test Dataset	Anticipative Calibration			Non-Anticipative Calibration			Regret (%)
	p^* (€/kWh)	$\frac{y^*-y}{\bar{y}-y}$ (%)	Profit (€)	p^* (€/kWh)	$\frac{y^*-y}{\bar{y}-y}$ (%)	Profit (€)	
2015	0.125	56.7	-74	n/a	n/a	n/a	n/a
2016	0.15		-101	0.15	56.7	-112	10
2017			-169			-169	0
2018			-190			-190	
2019	0.15	56.7	-153	0.15	56.7	-153	0

winning strategy. This strategy is non-anticipative if the year of the training dataset precedes the year of the test dataset. Table 1 shows that selecting p^* and y^* non-anticipatively on a historical training dataset has low regret relative to selecting these parameters anticipatively on the test dataset. Here, the regret is defined as the ratio of the absolute difference and the arithmetic mean of the cumulative profits generated by the anticipative and the non-anticipative strategies tuned on the test and training datasets, respectively. We always use the year immediately prior to the test dataset as training dataset. Table 1 shows that from 2017 onward, perhaps surprisingly, anticipative parameter tuning has no advantage over non-anticipative tuning. From now on, we assume that p^* and y^* are tuned non-anticipatively using the year of training data immediately prior to the test dataset. Additional robustness checks reveal that the cumulative profit (evaluated on 2019 data) is relatively insensitive to the choice of p^* and y^* within the suggested search grid, which indicates that its resolution is sufficiently high. Details are omitted for the sake of brevity.

6.3 Futility of Solving a Multi-Stage Model

The baseline strategy described in Section 6.2 looks only one day into the future. We now show numerically that the added value of solving a dynamic model that looks $H > 1$ days into the future is negligible. Specifically, we compare the baseline strategy against a dynamic strategy, which computes the market commitments on each day in a receding horizon fashion by solving model (5) with $H = 2$ and $\varphi_2(y_2) = p^*|y_2 - y^*|$, assuming that the electricity prices and the driving patterns for both subsequent days are known upfront. To compute the out-of-sample profits, we

use the backtesting procedure outlined in Section 6.2 but assume that the market bids are due at midnight when the initial state-of-charge y_0 is known. As explained in Section 5, this assumption greatly simplifies the computation of the dynamic strategy. We define the *value of vehicle-to-grid* under a particular bidding strategy as the cumulative *excess* profit of that strategy with respect to a simplified strategy that does not participate in the reserve market. This simplified strategy solves problem (R) under the additional constraint $x^r = 0$. When examining the value of vehicle-to-grid over the years from 2017 to 2019, we find that the dynamic strategy outperforms the baseline strategy only by 2.1%. Hence, the dynamic strategy does not generate significantly higher revenues even though it has the advantage of knowing the electricity prices and driving patterns *two* days in advance (which would not be the case in reality). We also emphasize that *both* strategies benefit from the assumption that y_0 is known upfront. However, we conjecture that both strategies benefit equally from this information relaxation and that the dynamic strategy thus still fails to outperform the baseline strategy when the market bids are due at noon and y_0 is uncertain. This reasoning justifies our use of a *static* baseline strategy that looks merely one day into the future.

6.4 Experiments: Set-up, Results and Discussion

In the remainder, we distinguish six different simulation scenarios. The nominal scenario uses the parameters of Table 2 in Appendix A for both training and testing. All other scenarios are based on slightly modified parameters. Specifically, we consider a lossless energy conversion scenario, which trains and tests the baseline strategy under the assumption that $\eta^+ = \eta^- = 1$. A variant of this scenario assumes lossless energy conversion in training but tests the resulting strategy under the nominal values of η^+ and η^- . We also consider two scenarios with weaker robustness guarantees that replace the uncertainty set \mathcal{D} in the training phase with its subset $\hat{\mathcal{D}}$. The resulting bidding strategy can be infeasible because it may fail to provide the legally required amount of reserve power. The two scenarios do not differ in training but impose different sanctions for infeasibilities in testing. In the first of the two scenarios the vehicle owner is immediately excluded from the reserve market upon the first infeasibility, thus losing the opportunity to earn money by offering grid services for the rest of the year. In the second scenario, the vehicle owner is penalized by $k_{\text{pen}} \cdot p^a(t)$ with $k_{\text{pen}} = 5$ for energy that is missing for frequency regulation (see also Section 6.2) and by $p^y = 0.75\text{€}/\text{kWh}$ for energy that is missing for driving. Finally, we consider a scenario in

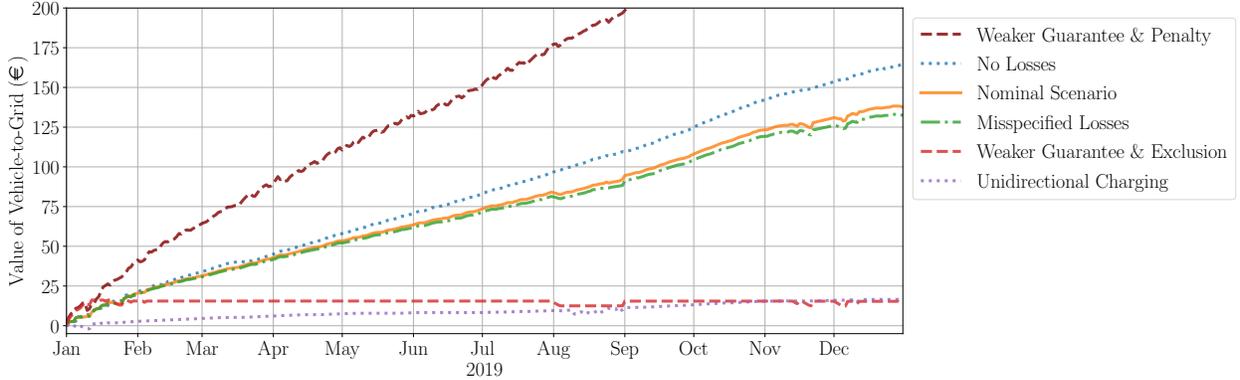


Figure 3. Value of vehicle-to-grid in 2019 in different simulation scenarios.

which the vehicle is only equipped with a unidirectional charger, that is, we set $\bar{y}^- = 0$. Thus, the vehicle is unable to feed power back into the grid. Requests for up-regulation ($\delta(t) < 0$) can therefore only be satisfied by consuming less energy, which is possible only if $\delta(t)x^r(t) \leq x^b(t)$.

Figure 3 visualizes the value of vehicle-to-grid in 2019 as a function of time for each of the six simulation scenarios. We first observe that the value of vehicle-to-grid in the lossless energy conversion scenario amounts to 165€ at the end of the year and thus significantly exceeds the respective value of 138€ in the nominal scenario. Using a perfectly efficient vehicle charger would thus have boosted the value of vehicle-to-grid by 20% in 2019. This is not surprising because a perfect charger prevents costly energy conversion losses. Note also that the scenario with misspecified efficiency parameters results in almost the same value of vehicle-to-grid as the nominal scenario, which suggests that misrepresenting η^+ or η^- in training has a negligible effect on the test performance. However, the underlying bidding strategy is *not* guaranteed to be feasible because it neglects energy conversion losses. Even though this strategy happens to remain feasible throughout 2019, it bears the risk of financial penalties or market exclusion. The two bidding strategies with weakened robustness guarantees initially reap high profits by aggressively participating in the reserve market, but they already fail in the first half of January to fulfill all market commitments. If infeasibilities are sanctioned by market exclusion, the cumulative excess profit thus remains flat after this incident. If infeasibilities lead to financial penalties, on the other hand, the cumulative excess profit drops sharply below zero near the incident but recovers quickly and then continues to grow steadily. As only a few other mild infeasibilities occur in 2019, the end-of-year excess profit of this aggressive bidding strategy still piles up to 296€, which is more than twice the excess profit in the nominal

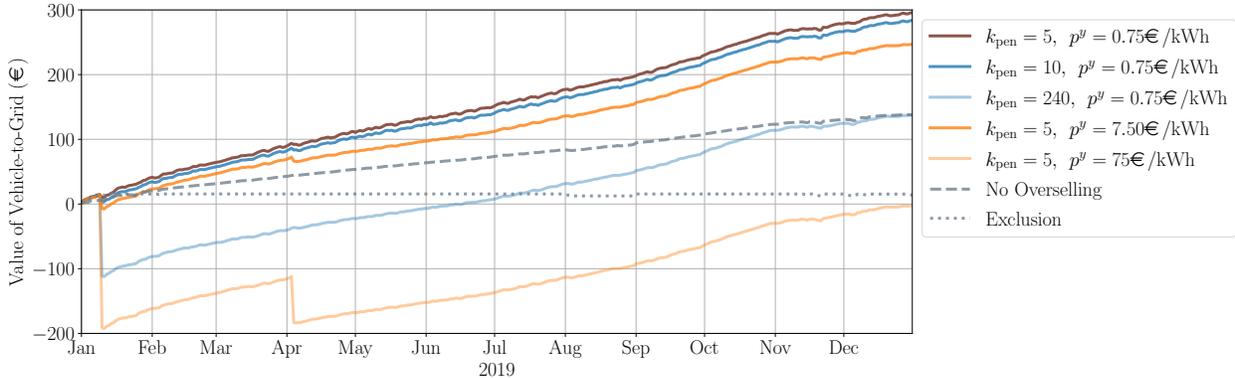


Figure 4. Value of vehicle-to-grid in 2019 for different penalty parameters.

scenario. We conclude that the current level of financial penalties is too low to deter vehicle owners from making promises they cannot honor. Finally, with a unidirectional charger, the 2019 value of vehicle-to-grid falls to 15€, which is only 11% of the respective value with a bidirectional charger.

As the bidding strategy with a weakened robustness guarantee can earn high profits when infringements of the EU delivery guarantee incur only financial penalties, we carry out an additional experiment to analyze the impact of the penalty parameters k_{pen} and p^y on the value of vehicle-to-grid; see Figure 4. We observe that for $p^y = 0.75\text{€}/\text{kWh}$, doubling the penalty factor to $k_{\text{pen}} = 10$ decreases the value of vehicle-to-grid by only 4.1% to 284€. An additional calculation reveals that the TSO would have to set $k_{\text{pen}} \approx 240$ in order to push the value of vehicle-to-grid below that attained in the baseline scenario, which fulfills the EU delivery guarantee. On the other hand, for $k_{\text{pen}} = 5$, a tenfold increase of p^y to $7.50\text{€}/\text{kWh}$ decreases the value of vehicle-to-grid by 16.6% to 247€, and an additional tenfold increase of p^y to $75\text{€}/\text{kWh}$ pushes the value of vehicle-to-grid below zero. Increasing p^y from $0.75\text{€}/\text{kWh}$ to $7.50\text{€}/\text{kWh}$ also reduces the number of days on which the vehicle owner does not have enough energy to drive from 7 to 2.

Since the market prices display distinct regime shifts as European energy policies evolve, the value of vehicle-to-grid fluctuates over the years. Indeed, Figure 8 in Appendix B shows that the value of vehicle-to-grid averages to 108€ per year from 2016 to 2019 with a minimum of 76€ in 2018.

In the next experiment we investigate how the value of vehicle-to-grid depends on the activation ratio γ/Γ and the battery’s size and charge rate (C-rate). The C-rate is defined as the ratio of the charger power and the battery size, and thus it expresses the percentage of the battery’s total capacity that can be charged within one hour. Figure 5 shows that the value of vehicle-to-grid

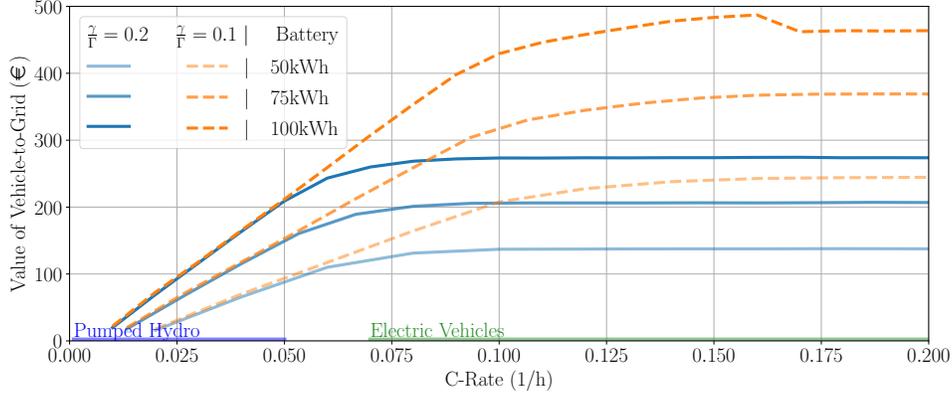


Figure 5. Value of vehicle-to-grid in 2019 as a function of the C-rate.

increases with the C-rate up to a saturation point that is insensitive to the battery size but decreases with the activation ratio. In the saturation regime, the value of vehicle-to-grid increases with the battery size. Typical electric vehicles can be fully charged overnight, within about 8 hours. The corresponding C-rate of 0.125h^{-1} falls within the saturation regime for both investigated activation ratios of 0.1 and 0.2. This observation has two implications. First, for typical electric vehicles the value of vehicle-to-grid cannot be increased by increasing the charger power (and thereby increasing the C-rate). This insight contradicts previous studies by Kempton and Tomić (2005), Codani et al. (2015) and Borne (2019), which advocate for electric vehicles with higher C-rates of 1, 0.45, and 0.37, respectively. The reason for this discrepancy is that none of the previous studies faithfully account for the EU delivery guarantee. In particular, Borne (2019) allows the vehicle owner to anticipate future frequency deviations, and Codani et al. (2015) assume that the bidding strategy can be updated on an hourly basis, thereby exploiting information that is not available at the time when the market bids are collected by the TSO (*i.e.*, one day in advance). By underestimating the amount of energy that the vehicle owner must be able to exchange with the TSO to satisfy all future obligations on the reserve market, these studies overestimate the amount of regulation power that can be sold, which makes potent battery chargers appear more useful than they actually are. The second implication is that the activation ratio has a critical impact on the value of vehicle-to-grid. The incumbent storage providers of the European electricity grid, namely pumped-hydro storage power plants, have C-rates of about 0.0125h^{-1} (Andrey et al., 2020), which are significantly smaller than the C-rates of electric vehicles. At such C-rates, the value of providing frequency

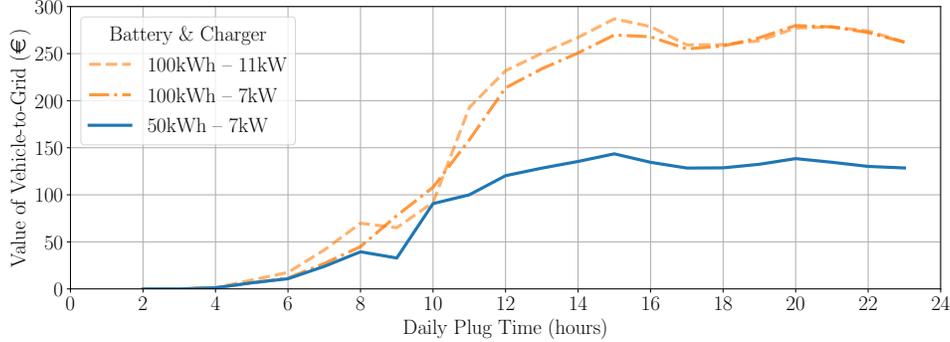


Figure 6. Value of vehicle-to-grid in 2019 versus daily plug time.

regulation is the same for activation ratios of 0.1 and 0.2. To minimize competition from vehicle-to-grid, pumped-hydro storage operators may thus have an incentive to lobby for high activation ratios. From the perspective of a TSO, the higher the activation ratio, the larger the uncertainty set \mathcal{D} and the lower the probability of blackouts. However, Figure 1 suggests that an activation period of 30 minutes is already conservative. On the other hand, the larger \mathcal{D} , the harder it is for storage operators to provide regulation power, which may lead to less competition, higher market prices, and a higher total cost of frequency regulation. As the system operator is a public entity, this cost is ultimately borne by the public, and the choice of the activation ratio is a political decision.

The last two experiments study the influence of the driving time and distance on the value of vehicle-to-grid. Figure 9 in Appendix B shows that the value of vehicle-to-grid in 2019 approximately displays a concave dependence on the yearly driving distance and plateaus at 138€ for driving distances between 5,000km and 10,000km. It is perhaps surprising that the value of vehicle-to-grid is *not* globally decreasing in the driving distance. An explanation for this phenomenon is provided in Appendix B. Figure 6 shows the value of vehicle-to-grid in 2019 against the daily plug time, that is, the total amount of time during which the vehicle is connected to a bidirectional charger. We observe that the value of vehicle-to-grid increases with the daily plug time and saturates after 15 hours at a level that scales with the battery size. Thus, a daily plug time of 15 hours suffices for offering the maximal possible amount of regulation power. Additional experiments show that for an activation ratio of 0.1 instead of 0.2, the saturation point increases to 20 hours.

7 Conclusions

We develop an optimization model for the decision problem of an electricity storage operator offering frequency regulation. To our best knowledge, this is the first model that faithfully accounts for the delivery guarantees required by the European Commission (2017). In contrast, all existing models relax the true delivery guarantee constraints and therefore risk that the electricity storage is empty (full) when a request for up- (down-)regulation arrives. In its original formulation, our model represents a non-convex robust optimization problem with functional uncertain parameters and thus appears to be severely intractable. Indeed, the state-of-the-art methods for solving the deterministic version of this problem for a single frequency deviation scenario in discrete time rely on techniques from mixed-integer linear programming. Maybe surprisingly, however, our robust optimization problem is equivalent to a tractable linear program. This is an exact result and does not rely on any approximations. To our best knowledge, we have thus discovered the first practically interesting class of optimization problems that become significantly easier through robustification.

In the numerical experiments centered around electric vehicles we restrict the planning horizon of our model to 24 hours. This is justified because the batteries of electric vehicles can be fully charged overnight and because vehicle owners may not know their driving needs several days in advance. Numerical experiments provide strong evidence that the added value of looking two or more days into the future is marginal even if future market prices and driving patterns were known.

As our optimization model faithfully captures effective legislation, it enables us to quantify the true value of vehicle-to-grid. This capability is relevant for understanding the economic incentives of different stakeholders such as vehicle owners, aggregators, equipment manufacturers, and regulators.

As for the vehicle owners, we find that their profits from frequency regulation range from 100€ to 500€ per year under typical driving patterns. It seems unlikely that such profits are sufficient for vehicle owners to forego the freedom of using their car whenever they please to. Nevertheless, some vehicle owners may choose to participate in vehicle-to-grid for idealistic reasons such as advancing the energy transition away from fossil fuels. Our numerical results also reveal that the value of vehicle-to-grid saturates at daily plug times above 15 hours. Thus, maximal profits from frequency regulation can be reaped even if the vehicle is disconnected from the grid up to 9 hours per day. This means that vehicle owners participating in vehicle-to-grid still enjoy considerable flexibility as to when to drive, which could help to promote the adoption of vehicle-to-grid.

Our results also have ramifications for aggregators, which pool multiple vehicles to offer regula-

tion power. Indeed, aggregators may allow vehicle owners to reserve their vehicles for up to 9 hours of driving per day without sacrificing profit. This leaves vehicle owners ample freedom and reduces the probability that they exceed their driving slots. Thus, the actual number of vehicles available for frequency regulation at any point in time closely matches its prediction. This allows aggregators to place reserve market bids with small safety margins, which ultimately lowers transaction costs.

Equipment manufacturers design and sell bidirectional vehicle chargers. Contrary to previous studies that relax the exact delivery guarantee constraints, we find that the battery size and not the charger power is limiting the profits from frequency regulation. Manufacturers thus have no incentive to produce overly powerful bidirectional chargers.

The electricity system and the society as a whole could benefit significantly from vehicle-to-grid, which harnesses the idle storage capacities of electric vehicles and thereby reduces the need for other sources of flexible electricity supply, such as gas power plants or stationary batteries. This in turn reduces the need for imports of natural gas and critical raw materials, increases the long-term security of electricity supply, and decreases greenhouse gas emissions. Regulators may therefore want to make vehicle-to-grid more attractive by prescribing the availability and delivery prices, defining appropriate delivery guarantee requirements and setting the penalties charged for non-compliance. Our results show that the vehicle owners' profits from frequency regulation decrease with the activation ratio and that current penalties are too low to incentivize vehicle owners to respect the law. Regulators could thus make vehicle-to-grid more attractive by decreasing the activation ratio and thereby relaxing the delivery guarantee requirements. Given that the delivery guarantee in our nominal scenario is very restrictive, this would only slightly decrease the reliability of frequency regulation. If, at the same time, regulators were to increase the penalties for non-compliance, then the reliability of frequency regulation might even increase because more vehicle owners would honor their contractual obligations. Our new model can be used for finding the appropriate level of the penalty. Our results also show that incumbent storage operators with low charge rates may lobby against weaker delivery guarantees because these would increase their competition from vehicle-to-grid but not increase their profits.

Acknowledgments D.L. acknowledges fruitful discussions with Nadège Faul, François Colet, Wilco Burghout, Paul Codani, Olivier Borne, Emilia Suomalainen, Jaâfar Berrada, Willett Kempton, Evangelos Vrettos, and Sophie Hiriart, as well as funding from the Institut VEDECOM.

References

- E. J. Anderson, P. Nash, and A. F. Perold. Some properties of a class of continuous linear programs. *SIAM Journal on Control and Optimization*, 21(5):758–765, 1983.
- C. Andrey, P. Barberi, L. Lacombe, L. van Nuffel, F. Gérard, J. Gorenstein Dedecca, K. Rademaekers, Y. El Idrissi, and M. Crenes. *Study on Energy Storage – Contribution to the Security of the Electricity Supply in Europe*. Publications Office of the European Union, 2020.
- E. Apostolaki-Iosifidou, P. Codani, and W. Kempton. Measurement of power loss during electric vehicle charging and discharging. *Energy*, 127:730–742, 2017.
- A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, Ser. A 99(2):351–376, 2004.
- A. Ben-Tal, L. El Ghaoui, and A. Nemirovski. *Robust Optimization*. Princeton University Press, 2009.
- D. Bertsimas and M. Sim. The price of robustness. *Operations Research*, 52(1):35–53, 2004.
- D. Bertsimas and J. N. Tsitsiklis. *Introduction to Linear Optimization*. Athena Scientific & Dynamic Ideas, 1997.
- O. Borne. *Vehicle-to-Grid and Flexibility for Electricity Systems: From Technical Solutions to Design of Business Models*. PhD thesis, Université Paris-Saclay, 2019.
- O. Borne, Y. Perez, and M. Petit. Market integration or bids granularity to enhance flexibility provision by batteries of electric vehicles. *Energy Policy*, 119:140–148, 2018.
- S. P. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- G. Broneske and D. Wozabal. How do contract parameters influence the economics of vehicle-to-grid? *Manufacturing & Service Operations Management*, 19(1):150–164, 2017.
- P. Carpentier, J.-P. Chancelier, M. de Lara, and T. Rigaut. Algorithms for two-time scales stochastic optimization with applications to long term management of energy storage. Preprint available from <http://hal.archives-ouvertes.fr/hal-02013963>, 2019.
- P. Codani, M. Petit, and Y. Perez. Participation of an electric vehicle fleet to primary frequency control in France. *International Journal of Electric and Hybrid Vehicles*, 7(3):233–249, 2015.
- Commissariat général au développement durable. La mobilité des Français. *La Revue du CGDD*, 2010.
- J. Donadee and M. D. Ilić. Stochastic optimization of grid to vehicle frequency regulation capacity bids. *IEEE Transactions on Smart Grid*, 5(2), 2014.
- M. Dubarry, A. Devie, and K. McKenzie. Durability and reliability of electric vehicle batteries under electric utility grid operations: Bidirectional charging impact analysis. *Journal of Power Sources*, 358:39–49, 2017.
- European Commission. Commission Regulation (EU) 2017/1485 of 2 August 2017 establishing a guideline on electricity transmission system operation. *Official Journal of the European Union*, 60, 2017.
- Federal Highway Administration. 2017 National Household Travel Survey. Technical report, U.S. Department of Transportation, 2017. URL <http://nhts.ornl.gov>.
- A. Georghiou, A. Tsoukalas, and W. Wiesemann. Robust dual dynamic programming. *Operations Research*, 67(3):813–830, 2019.

- Gestionnaire du Réseau de Transport d'Electricité. Documentation Technique de Référence, Chapitre 4–Contribution des utilisateurs aux performances du RPT, Article 4.1–Réglage Fréquence/Puissance. Technical report, 2009.
- A. Ghate. Robust continuous linear programs. *Optimization Letters*, 14:1627–1642, 2020.
- J. D. Glover, M. S. Sarma, and T. J. Overbye. *Power Systems Analysis and Design*. Cengage Learning, 2010.
- C. Guille and G. Gross. A conceptual framework for the *vehicle-to-grid (V2G)* implementation. *Energy Policy*, 37:4379–4390, 2009.
- S. Han, S. Han, and K. Sezaki. Development of an optimal vehicle-to-grid aggregator for frequency regulation. *IEEE Transactions on Smart Grid*, 1(1):65–72, 2010.
- S. Han, S. Han, and K. Sezaki. Estimation of achievable power capacity from plug-in electric vehicles for V2G frequency regulation: Case studies for market participation. *IEEE Transactions on Smart Grid*, 2:632–641, 2011.
- G. He, Q. Chen, C. Kang, P. Pinson, and Q. Xia. Optimal bidding strategy of battery storage in power markets considering performance-based regulation and battery cycle life. *IEEE Transactions on Smart Grid*, 7(5):2359–2367, 2016.
- B. Houska. *Robust Optimization of Dynamic Systems*. PhD thesis, KU Leuven, 2011.
- W. Kempton and S. E. Letendre. Electric vehicles as a new power source for electric utilities. *Transportation Research Part D: Transport and Environment*, 2(3):157–175, 1997.
- W. Kempton and J. Tomić. Vehicle-to-grid power fundamentals: Calculating capacity and net revenue. *Journal of Power Sources*, 144(1):268–279, 2005.
- D. Lauinger, F. Vuille, and D. Kuhn. A review of the state of research on vehicle-to-grid (V2G): Progress and barriers to deployment. In *European Battery, Hybrid and Fuel Cell Electric Vehicle Congress*, Geneva, 2017.
- N. Löhdorf, D. Wozabal, and S. Minner. Optimizing trading decisions for hydro storage systems using approximate dual dynamic programming. *Operations Research*, 61(4):810–823, 2013.
- P. Mercier, R. Cherkaoui, and A. Oudalov. Optimizing a battery energy storage system for frequency control application in an isolated power system. *IEEE Transactions on Power Systems*, 24(3):1469–1477, 2009.
- S. Nadarajah, F. Margot, and N. Secomandi. Relaxations of approximate linear programs for the real option management of commodity storage. *Management Science*, 61(12):3054–3076, 2015.
- E. Namor, F. Sossan, R. Cherkaoui, and M. Paolone. Control of battery storage systems for the simultaneous provision of multiple services. *IEEE Transactions on Smart Grid*, 10(3):2799–2808, 2019.
- G. L. Nemhauser and L. A. Wolsey. *Integer and Combinatorial Optimization Part III*. John Wiley & Sons, 1999.
- L. Noel, G. Z. de Rubens, J. Kester, and B. K. Sovacool. *Vehicle-to-Grid: A Sociotechnical Transition Beyond Electric Mobility*. Palgrave Macmillan, 2019.
- M. V. F. Pereira and L. M. V. G. Pinto. Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming*, 52(2):359–375, 1991.

- G. Pritchard, A. Philpott, and P. Neame. Hydroelectric reservoir optimization in a pool market. *Mathematical Programming*, Ser. A 103(3):445–461, 2005.
- M. C. Pullan. Forms of optimal solutions for separated continuous linear programs. *SIAM Journal on Control and Optimization*, 33(6):1952–1977, 1995.
- Y. G. Rebours, D. S. Kirschen, M. Trotignon, and S. Rossignol. A survey of frequency and voltage control ancillary services—Part I: Technical features. *IEEE Transactions on Power Systems*, 22(1):350–357, 2007.
- Réseau de transport d’électricité. Règles Service Système Fréquence. Technical report, 2017a.
- Réseau de transport d’électricité. Bilan prévisionnel de l’équilibre offre-demande d’électricité en France. Technical report, 2017b.
- Réseau de transport d’électricité. Electricity Report 2017. Technical report, 2017c.
- E. Roos, D. den Hertog, A. Ben-Tal, F. de Ruiter, and J. Zhen. Tractable approximation of hard uncertain convex optimization problems. *Available from Optimization Online*, 2020.
- N. Secomandi. Optimal commodity trading with a capacitated storage asset. *Management Science*, 56(3):449–467, 2010.
- E. Sortomme and M. A. El-Sharkawi. Optimal scheduling of vehicle-to-grid energy and ancillary services. *IEEE Transactions on Smart Grid*, 3(1):351–359, 2012.
- D. Sperling. *Future Drive: Electric Vehicles and Sustainable Transportation*. Island Press, Washington, D.C. & Covelo, California, 1994.
- N. Sunar and J. R. Birge. Strategic commitment to a production schedule with uncertain supply and demand: Renewable energy in day-ahead electricity markets. *Management Science*, 65(2):714–734, 2019.
- T. M. Sweda, I. S. Dolinskaya, and D. Klabjan. Optimal recharging policies for electric vehicles. *Transportation Science*, 51(2):457–479, 2017.
- J. A. Taylor. *Convex Optimization of Power Systems*. Cambridge University Press, 2015.
- A. Thompson. Economic implications of lithium ion battery degradation for vehicle-to-grid (V2X) services. *Journal of Power Sources*, 396:691–709, 2018.
- K. Uddin, T. Jackson, W. D. Widanage, G. Chouchelamane, P. A. Jennings, and J. Marco. On the possibility of extending the lifetime of lithium-ion batteries through optimal V2G facilitated by an integrated vehicle and smart-grid system. *Energy*, 133:710–722, 2017.
- K. Uddin, M. Dubarry, and M. B. Glick. The viability of vehicle-to-grid operations from a battery technology and policy perspective. *Energy Policy*, 113:342–347, 2018.
- S. Vandael, T. Holvoet, G. Deconinck, S. Kamboj, and W. Kempton. A comparison of two GIV mechanisms for providing ancillary services at the University of Delaware. In *IEEE SmartGridComm 2013 Symposium—Demand Side Management, Demand Response, Dynamic Pricing*, 2013.
- S. Vandael, T. Holvoet, G. Deconinck, H. Nakano, and W. Kempton. A scalable control approach for providing regulation services with grid-integrated electric vehicles. In T. Suzuki, S. Inagaki, Y. Susuki, and A. T. Tran, editors, *Design and Analysis of Distributed Energy Management Systems*, chapter 6, pages 107–128. Springer, 2020.
- J. Warrington, P. Goulart, S. Mariethoz, and M. Morari. Policy-based reserves for power systems. *IEEE Transactions on Power Systems*, 20(4):4427–4437, 2013.

- G. Wenzel, M. Negrete-Pincetic, D. E. Olivares, J. MacDonald, and D. S. Callaway. Real-time charging strategies for an electric vehicle aggregator to provide ancillary services. *IEEE Transactions on Smart Grid*, 9(5):5141–5151, 2018.
- F. Wu and R. Sioshansi. A stochastic operational model for controlling electric vehicle charging to provide frequency regulation. *Transportation Research Part D*, 67:475–490, 2019.
- E. Yao, V. W. S. Wong, and R. Schober. Robust frequency regulation capacity scheduling algorithm for electric vehicles. *IEEE Transactions on Smart Grid*, 8(2):984–997, 2017.
- W. W.-G. Yeh. Reservoir management and operations models: A state-of-the-art review. *Water Resources Research*, 21(12):1797–1818, 1985.
- X. Zhang, M. Kamgarpour, A. Georghiou, P. Goulart, and J. Lygeros. Robust optimal control with adjustable uncertainty sets. *Automatica*, 75:249–259, 2017.
- Y. Zhang, M. Lu, and S. Shen. On the values of vehicle-to-grid electricity selling in electric vehicle sharing. *Manufacturing & Service Operations Management*, 23(2):488–507, 2021.
- Y. H. Zhou, A. Scheller-Wolf, N. Secomandi, and S. Smith. Electricity trading and negative prices: Storage vs. disposal. *Management Science*, 62(3):880–898, 2016.
- S. Zymler, B. Rustem, and D. Kuhn. Robust portfolio optimization with derivative insurance guarantees. *European Journal of Operational Research*, 210(2):410–424, 2011.

A Problem Data and Model Parameters

Table 2. Parameters of the Nominal Simulation Scenario.

Parameter	Symbol	Value	Unit
Vehicle Data			
Minimum State-of-Charge	\underline{y}	10	kWh
Maximum State-of-Charge	\bar{y}	40	kWh
Target State-of-Charge	y^*	27	kWh
Deviation Penalty	p^*	0.15	€/kWh
Charging Efficiency	η^+	85	%
Discharging Efficiency	η^-	85	%
Maximum Charging Power	\bar{y}^+	7	kW
Maximum Discharging Power	\bar{y}^-	7	kW
Yearly Energy for Driving		2,000	kWh
Fraction of Time Driving		27	%
Grid Data			
Nominal frequency	f_0	50	Hz
Normalization constant	Δf	200	mHz
Average Utility Price from 2015 to 2019		14.31	cts/kWh
Average Availability Price from 2015 to 2019		0.825	cts/kW/h
General Parameters			
Trading Interval	Δt	30	min
Activation Period in \mathcal{D}	γ	30	min
Regulation Cycle in \mathcal{D}	Γ	2.5	h
Activation Period in $\hat{\mathcal{D}}$	$\hat{\gamma}$	30	min
Regulation Cycle in $\hat{\mathcal{D}}$	$\hat{\Gamma}$	1	day
Sampling Rate for Frequency Measurements		0.1	Hz
Planning Horizon	T	1	day



Figure 7. Evolution of the availability and regulation prices from 2015 to 2019. The regulation price changes every 10s. For better visibility, we show its daily averages.

B Additional Experiments and Results

In the first additional experiment we assess the heterogeneity of the value of vehicle-to-grid across the years 2016 to 2019. Figure 8 visualizes the temporal evolution of the value of vehicle-to-grid and shows that it may vary between 75€ and 140€ at the end of the year. The second additional experiment compares two vehicles with uni- and bidirectional chargers. Figure 9 indicates that the value of vehicle-to-grid is approximately concave in the yearly mileage for both vehicles and for mileages up to 30,000km. Since a vehicle with a bidirectional charger can be used for unidirectional charging, the value of vehicle-to-grid with a bidirectional charger exceeds that of the same vehicle with a unidirectional charger. Furthermore, a higher mileage necessitates higher utility purchases x^b , which has two opposing effects on regulation profits. On the one hand, power discharges become less likely because they only occur when $\delta(t)x^r(t) < -x^b(t)$, which reduces energy conversion losses and makes the provision of frequency regulation more cost-effective. On the other hand, the effective upper bound $\bar{y}^+(t) - x^b(t)$ on $x^r(t)$ tightens, which reduces the amount of regulation power that can be offered on the market. The value of vehicle-to-grid peaks at 138€ for a vehicle with a bidirectional charger traveling between 5,000km and 10,000km per year and at 42€ for a vehicle with a unidirectional charger traveling 30,000km per year. For yearly mileages greater than 5,000km, the higher the yearly mileage, the lower the added value of a bidirectional charger. This result may be of particular interest for operators of shared electric vehicles. However, even for a yearly mileage of 30,000km, using a bidirectional instead of a unidirectional charger more than doubles the value of vehicle-to-grid. Furthermore, many drivers may prefer vehicles with internal combustion engines or fuel cells over electric vehicles to cover high yearly mileages because of their greater

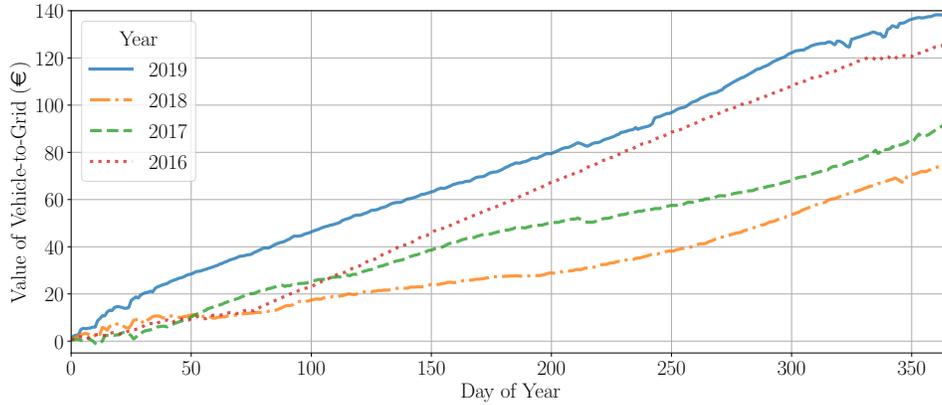


Figure 8. Value of vehicle-to-grid from 2016 to 2019.

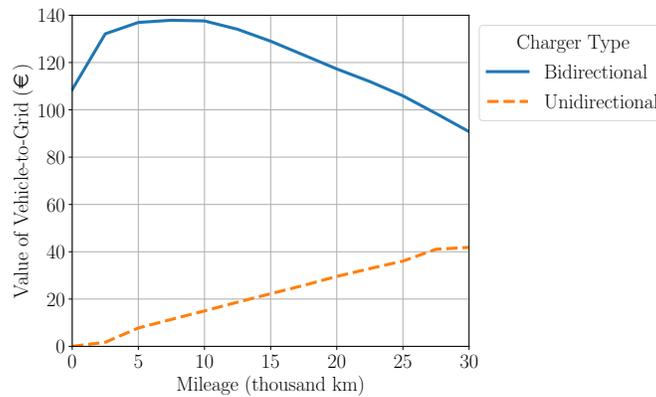


Figure 9. Value of vehicle-to-grid in 2019 vs mileage.

ranges and shorter refueling times. Therefore, owners of electric vehicle are more likely to cover yearly mileages close to the French average of 13,000km. Providing frequency regulation with a unidirectional charger would thus earn them about 20€ per year. In practice, these earnings would have to be shared with vehicle aggregators and equipment manufacturers.

To complement the analysis of the nominal scenario described in Section 6.4, Figure 10 shows the relationships between market bids, prices, driving needs, frequency deviations, and the battery state-of-charge on 9 August 2019 in a 10 second resolution. As expected, the vehicle owner charges the battery at night when utility prices are low. We also observe that the vehicle participates in the regulation market at a more or less constant level whenever it is not driving. This reduces the exposure to extreme frequency deviations at any one time. The battery state-of-charge naturally decreases when the vehicle is driving, increases when the vehicle is being charged and remains

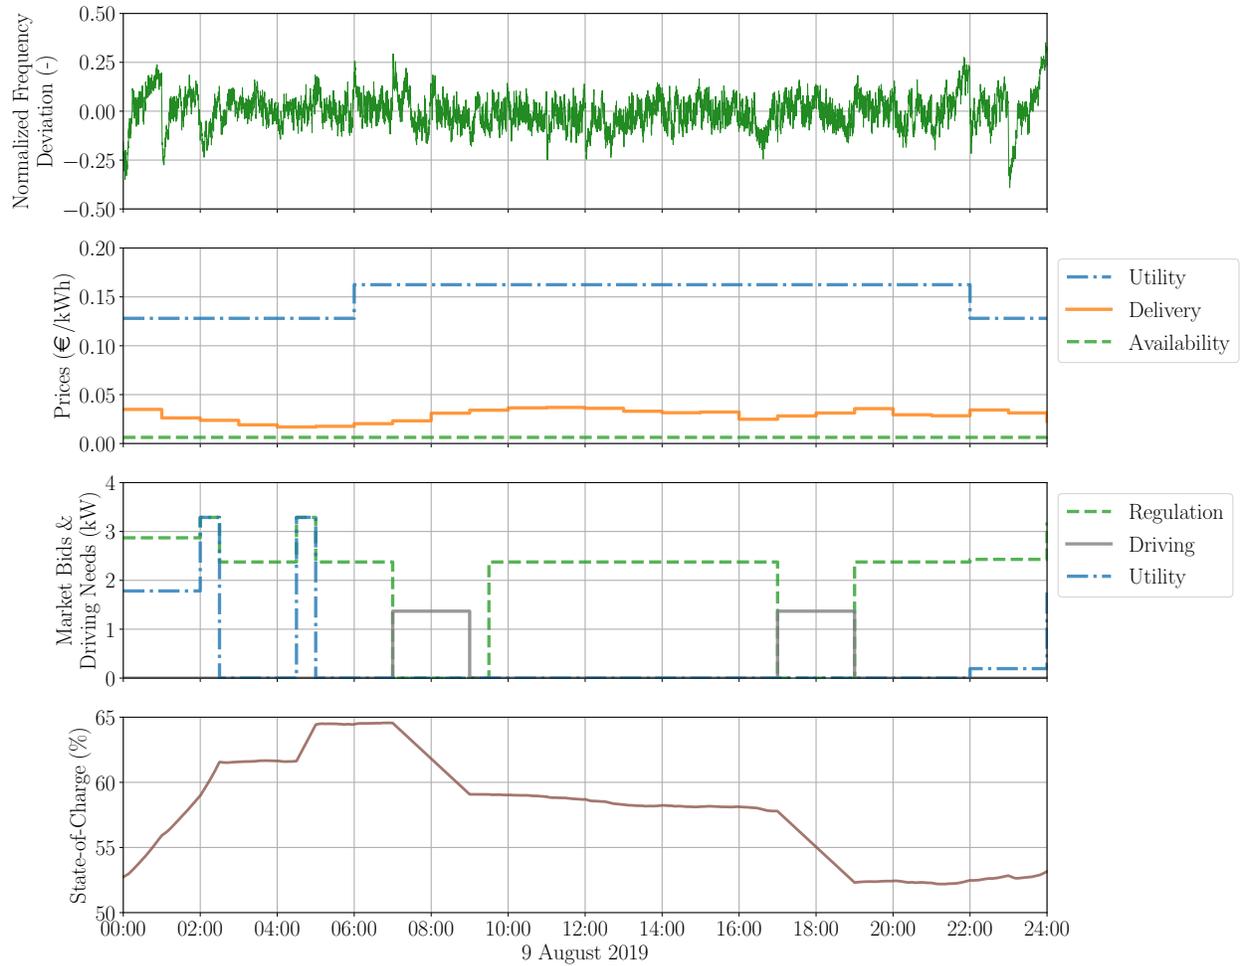


Figure 10. Frequency deviations, prices, market bids, driving needs, and state-of-charge on 9 August 2019.

essentially constant when the vehicle provides only frequency regulation because the frequency deviations fluctuate rapidly around zero.

C Proofs

This appendix contains the proofs of all theorems and propositions in the main text.

Proof of Proposition 1. By definition we have

$$\begin{aligned} y(x^b, x^r, \delta, y_0, t) &= y_0 + \int_0^t \eta^+ [x^b(t') + \delta(t')x^r(t')]^+ - \frac{[x^b(t') + \delta(t')x^r(t')]^-}{\eta^-} - d(t') dt' \\ &= y_0 + \int_0^t \min \left\{ \eta^+ (x^b(t') + \delta(t')x^r(t')), \frac{x^b(t') + \delta(t')x^r(t')}{\eta^-} \right\} - d(t') dt', \end{aligned}$$

where the second equality holds because $\eta^+ < \frac{1}{\eta^-}$. The postulated properties of $y(x^b, x^r, \delta, y_0, t)$ follow from the observation that the minimum of two (nondecreasing) affine functions is a concave (nondecreasing) function (Boyd and Vandenberghe, 2004, p. 73). \square

Proof of Proposition 2. We will show that $\max_{\delta \in \mathcal{D}} \text{Var}(\delta) \leq \beta/T$, where $\beta = \lceil T/\Gamma \rceil \gamma$. To this end, we note that

$$\max_{\delta \in \mathcal{D}} \text{Var}(\delta) = \max_{\delta \in \mathcal{D}^+} \text{Var}(\delta) \leq \max_{\delta \in \mathcal{D}_\beta} \text{Var}(\delta), \quad (6)$$

where $\mathcal{D}_\beta = \{\delta \in \mathcal{L}(\mathcal{T}, [0, 1]) : \int_0^T \delta(t) dt \leq \beta\}$. The equality in (6) holds because $\text{Var}(\delta)$ remains unchanged when $\delta(t)$ is replaced with $|\delta(t)|$ for every $t \in \mathcal{T}$. Moreover, the inequality holds because $\mathcal{D}^+ \subseteq \mathcal{D}_\beta$. To see this, observe that for any $\delta \in \mathcal{D}^+$ we have

$$\int_0^T \delta(t) dt = \sum_{n=1}^{\lceil T/\Gamma \rceil} \int_{(n-1)\Gamma}^{\min\{n\Gamma, T\}} \delta(t) dt \leq \left\lceil \frac{T}{\Gamma} \right\rceil \gamma = \beta.$$

Note that \mathcal{D}_β constitutes a budget uncertainty set of the type introduced by Bertsimas and Sim (2004), where the uncertainty budget β corresponds to the maximum number of regulation cycles within the planning horizon multiplied by the duration of an activation period. Thus, β can be viewed as the maximum amount of time within the planning horizon during which all reserve commitments must be honored. By weak duality, the highest variance of any scenario in \mathcal{D}_β satisfies

$$\begin{aligned} \max_{\delta \in \mathcal{D}_\beta} \text{Var}(\delta) &\leq \min_{\lambda \geq 0} \max_{\delta \in \mathcal{L}(\mathcal{T}, [0, 1])} \frac{1}{T} \int_0^T \delta^2(t) dt + \lambda \left(\beta - \int_0^T \delta(t) dt \right) \\ &= \min_{\lambda \geq 0} \lambda \beta + \max_{\delta \in \mathcal{L}(\mathcal{T}, [0, 1])} \int_0^T \delta(t) (\delta(t)/T - \lambda) dt \\ &= \min_{\lambda \geq 0} \lambda \beta + T \max_{\delta \in [0, 1]} \delta (\delta/T - \lambda) \\ &= \min_{\lambda \geq 0} \lambda \beta + \max \{0, 1 - \lambda T\} \\ &= \min_{\lambda \geq 0} \max \{\lambda \beta, 1 - \lambda(T - \beta)\} = \frac{\beta}{T}. \end{aligned}$$

The claim now follows by substituting the above result into (6) and recalling the definition of β . \square

Proof of Proposition 3. As for assertion (i) we first prove that $L\mathcal{D}_{\mathcal{K}} \subseteq \mathcal{D}$. To this end, select any $\delta \in \mathcal{D}_{\mathcal{K}}$, and define $\delta = L\delta$. It is easy to see that $\delta \in \mathcal{L}(\mathcal{T}, [-1, 1])$. Next, select any $t \in \mathcal{T}$. If $t \leq \Gamma$, note that

$$\int_{[t-\Gamma]^+}^t |\delta(t')| dt' = \int_0^t |\delta(t')| dt' \leq \int_0^\Gamma |\delta(t')| dt' = \Delta t \sum_{l=1}^{\Delta k} |\delta_l| \leq \gamma,$$

where the auxiliary parameter $\Delta k = \Gamma/\Delta t$ is integral thanks to Assumption 2. The second inequality in the above expression holds because $\delta \in \mathcal{D}_{\mathcal{K}}$. If $t \geq \Gamma$, on the other hand, we define $k = \lceil \frac{t}{\Delta t} \rceil$ and $\alpha = \lceil \frac{t}{\Delta t} \rceil - \frac{t}{\Delta t} \in [0, 1)$. Then, we find

$$\begin{aligned} \int_{[t-\Gamma]^+}^t |\delta(t')| dt' &= \int_{t-\Gamma}^{(k-\Delta k)\Delta t} |\delta(t')| dt' + \int_{(k-\Delta k)\Delta t}^{(k-1)\Delta t} |\delta(t')| dt' + \int_{(k-1)\Delta t}^t |\delta(t')| dt' \\ &= (k\Delta t - t) |\delta_{k-\Delta k}| + \Delta t \sum_{l=k-\Delta k+1}^{k-1} |\delta_l| + (t - (k-1)\Delta t) |\delta_k| \\ &= \Delta t \left(\alpha \sum_{l=k-\Delta k}^{k-1} |\delta_l| + (1-\alpha) \sum_{l=k-\Delta k+1}^k |\delta_l| \right) \leq \gamma, \end{aligned}$$

where the inequality holds because $\delta \in \mathcal{D}_{\mathcal{K}}$, which ensures that both $\sum_{l=k-\Delta k}^{k-1} |\delta_l|$ and $\sum_{l=k-\Delta k+1}^k |\delta_l|$ are smaller or equal to $\gamma/\Delta t$. As $t \in \mathcal{T}$ was chosen arbitrarily, this implies that $\delta \in L\mathcal{D}_{\mathcal{K}}$. In summary, we have shown that $L\mathcal{D}_{\mathcal{K}} \subseteq \mathcal{D}$.

Next, we show that $L^\dagger \mathcal{D} \subseteq \mathcal{D}_{\mathcal{K}}$. To this end, select any $\delta \in \mathcal{D}$ and define $\delta = L^\dagger \delta$. It is easy to see that $\delta \in [-1, 1]^K$. Moreover, for any $k \in \mathcal{K}$ we have

$$\sum_{l=1+[k-\Gamma/\Delta t]^+}^k \delta_l = \sum_{l=1+[k-\Gamma/\Delta t]^+}^k \frac{1}{\Delta t} \int_{\mathcal{T}_l} \delta(t') dt' = \frac{1}{\Delta t} \int_{[k\Delta t-\Gamma]^+}^{k\Delta t} \delta(t') dt' \leq \frac{\gamma}{\Delta t},$$

where the inequality holds because $\delta \in \mathcal{D}$. As $k \in \mathcal{K}$ was chosen arbitrarily, this implies that $\delta \in L^\dagger \mathcal{D}$. In summary, we have shown that $L^\dagger \mathcal{D} \subseteq \mathcal{D}_{\mathcal{K}}$.

Finally, we prove that $\mathcal{D}_{\mathcal{K}} \subseteq L^\dagger \mathcal{D}$. To this end, we observe that $L^\dagger L$ coincides with the identity mapping on \mathbb{R}^K . As $L\mathcal{D}_{\mathcal{K}} \subseteq \mathcal{D}$, this implies that

$$\mathcal{D}_{\mathcal{K}} = L^\dagger L\mathcal{D}_{\mathcal{K}} \subseteq L^\dagger \mathcal{D}.$$

Since both $L^\dagger \mathcal{D} \subseteq \mathcal{D}_{\mathcal{K}}$ and $\mathcal{D}_{\mathcal{K}} \subseteq L^\dagger \mathcal{D}$, we have in fact shown that $L^\dagger \mathcal{D} = \mathcal{D}_{\mathcal{K}}$. Thus assertion (i) follows. Assertion (ii) can be proved in a similar manner. Details are omitted for brevity. \square

Proof of Proposition 4. The proof widely parallels that of Proposition 1 and is therefore omitted. \square

The proof of Theorem 1 relies on Propositions 6–9 below.

Proposition 6. *The following equivalences hold.*

$$(i) \quad y^+ \left(x^b(t), x^r(t), \delta(t) \right) \leq \bar{y}^+(t) \quad \forall \delta \in \mathcal{D}, \forall t \in \mathcal{T} \iff y_k^+ \left(x_k^b, x_k^r, \delta_k \right) \leq \bar{y}_k^+ \quad \forall \delta \in \mathcal{D}_{\mathcal{K}}, \forall k \in \mathcal{K}$$

$$(ii) \quad y^- \left(x^b(t), x^r(t), \delta(t) \right) \leq \bar{y}^-(t) \quad \forall \delta \in \mathcal{D}, \forall t \in \mathcal{T} \iff y_k^- \left(x_k^b, x_k^r, \delta_k \right) \leq \bar{y}_k^- \quad \forall \delta \in \mathcal{D}_{\mathcal{K}}, \forall k \in \mathcal{K}$$

Proof. Assertion (i) can be reexpressed as

$$\max_{t \in \mathcal{T}} \max_{\delta \in \mathcal{D}} y^+ \left(x^b(t), x^r(t), \delta(t) \right) - \bar{y}^+(t) \leq 0 \iff \max_{k \in \mathcal{K}} \max_{\delta \in \mathcal{D}_{\mathcal{K}}} y_k^+ \left(x_k^b, x_k^r, \delta_k \right) - \bar{y}_k^+ \leq 0.$$

We will prove this equivalence by showing that the left hand sides of the two inequalities are equal. Indeed, a direct calculation reveals that

$$\begin{aligned} \max_{t \in \mathcal{T}} \max_{\delta \in \mathcal{D}} y^+ \left(x^b(t), x^r(t), \delta(t) \right) - \bar{y}^+(t) &= \max_{t \in \mathcal{T}} \max_{-1 \leq \delta(t) \leq 1} \left[x^b(t) + \delta(t)x^r(t) \right]^+ - \bar{y}^+(t) \\ &= \max_{t \in \mathcal{T}} x^b(t) + x^r(t) - \bar{y}^+(t) \\ &= \max_{k \in \mathcal{K}} x_k^b + x_k^r - \bar{y}_k^+ \\ &= \max_{k \in \mathcal{K}} \max_{-1 \leq \delta_k \leq 1} \left[x_k^b + \delta_k x_k^r \right]^+ - \bar{y}_k^+ \\ &= \max_{k \in \mathcal{K}} \max_{\delta \in \mathcal{D}_{\mathcal{K}}} y^+ \left(x_k^b, x_k^r, \delta_k \right) - \bar{y}_k^+, \end{aligned} \tag{8}$$

where the first equality follows from the definition of y^+ in (1a) and the observation that $\{\delta(t) : \delta \in \mathcal{D}\} = [-1, 1]$, while the second equality holds because $x^b(t) \geq 0$ and $x^r(t) \geq 0$ which implies that $\delta(t) = 1$ maximizes the instantaneous charging rate. The third equality exploits our assumption that x^b, x^r , and \bar{y}^+ are piecewise constant functions. The fourth equality holds because $x_k^b \geq 0$ and $x_k^r \geq 0$, which implies that $\delta_k = 1$ maximizes the per-period charging rate. The fifth equality follows again from the definition of y^+ in (1a) and the observation that $\{\delta_k : \delta \in \mathcal{D}_{\mathcal{K}}\} = [-1, 1]$.

The proof of assertion (ii) is similar and therefore omitted. \square

Proposition 7. *The following equivalence holds.*

$$y \left(x^b, x^r, \delta, y_0, t \right) \leq \bar{y} \quad \forall \delta \in \mathcal{D}, \forall t \in \mathcal{T} \iff y_k \left(\mathbf{x}^b, \mathbf{x}^r, \boldsymbol{\delta}, y_0 \right) \leq \bar{y} \quad \forall \delta \in \mathcal{D}_{\mathcal{K}}, \forall k \in \mathcal{K} \cup \{0\}$$

Proof. The claim follows if we can show that

$$\max_{t \in \mathcal{T}} \max_{\delta \in \mathcal{D}} y \left(x^b, x^r, \delta, y_0, t \right) = \max_{k \in \mathcal{K} \cup \{0\}} \max_{\delta \in \mathcal{D}_{\mathcal{K}}} y_k \left(\mathbf{x}^b, \mathbf{x}^r, \boldsymbol{\delta}, y_0 \right). \tag{9}$$

To this end, assume first that $t = k\Delta t$ for some $k \in \mathcal{K} \cup \{0\}$. In this case, we have

$$\begin{aligned}
\max_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, t) &= \max_{\delta \in \mathcal{D}^+} y(x^b, x^r, \delta, y_0, t) \\
&= y_0 + \max_{\delta \in \mathcal{D}^+} \int_0^t \eta^+ \left(x^b(t') + \delta(t')x^r(t') \right) - d(t') dt' \\
&= y_0 + \max_{\delta \in \mathcal{D}^+} \sum_{l=1}^k \int_{\mathcal{T}_l} \eta^+ \left(x_l^b + \delta(t')x_l^r \right) - d_l dt' \\
&= y_0 + \max_{\delta \in \mathcal{D}^+} \Delta t \sum_{l=1}^k \eta^+ \left(x_l^b + (L^\dagger \delta)_l x_l^r \right) - d_l \\
&= y_0 + \max_{\delta \in \mathcal{D}_{\mathcal{K}}^+} \Delta t \sum_{l=1}^k \eta^+ \left(x_l^b + \delta_l x_l^r \right) - d_l \\
&= \max_{\delta \in \mathcal{D}_{\mathcal{K}}^+} y_k(x^b, x^r, \delta, y_0) = \max_{\delta \in \mathcal{D}_{\mathcal{K}}} y_k(x^b, x^r, \delta, y_0),
\end{aligned} \tag{10}$$

where the first equality holds because $\delta \in \mathcal{D}$ if and only if $|\delta| \in \mathcal{D}^+$ and because y is nondecreasing in δ thanks to Proposition 1. The second equality follows from the definitions of y , y^+ , and y^- and from the non-negativity of x^b , x^r and δ . The third equality exploits our assumption that d , x^b and x^r are piecewise constant. As δ is integrated against a piecewise constant function, it may be averaged over the trading intervals without changing its objective function value. The fifth equality then follows from Proposition 3, while the sixth equality follows from the definitions of y_k , y^+ and y^- and from the non-negativity of x^b , x^r and δ . The seventh equality, finally, holds because $\delta \in \mathcal{D}_{\mathcal{K}}$ if and only if $|\delta| \in \mathcal{D}_{\mathcal{K}}^+$ and because y_k is nondecreasing in δ thanks to Proposition 6.

Assume now more generally that $t \in \mathcal{T}_k$ for some $k \in \mathcal{K}$. If the vehicle is driving in trading interval \mathcal{T}_k , then $\bar{y}^+(t) = \bar{y}^-(t) = 0$ for all $t \in \mathcal{T}_k$. Thus, we have

$$\begin{aligned}
\max_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, t) &= \max_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, (k-1)\Delta t) - \int_{(k-1)\Delta t}^t d(t') dt' \\
&\leq \max_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, (k-1)\Delta t) = \max_{\delta \in \mathcal{D}_{\mathcal{K}}} y_{k-1}(x^b, x^r, \delta, y_0) \quad \forall t \in \mathcal{T}_k,
\end{aligned}$$

where the inequality holds because $d(t) \geq 0$ for all $t \in \mathcal{T}_k$, and the second equality follows from the first part of the proof. Alternatively, if the vehicle is parked in trading interval \mathcal{T}_k , then $d(t) = 0$ for all $t \in \mathcal{T}_k$. Thus, we have

$$\begin{aligned}
\max_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, t) &= \max_{\delta \in \mathcal{D}^+} y(x^b, x^r, \delta, y_0, t) \\
&= \max_{\delta \in \mathcal{D}^+} y(x^b, x^r, \delta, y_0, (k-1)\Delta t) + \int_{(k-1)\Delta t}^t \eta^+ y^+ \left(x^b(t'), x^r(t'), \delta(t') \right) dt'
\end{aligned}$$

$$\leq \max_{\delta \in \mathcal{D}^+} y(x^b, x^r, \delta, y_0, k\Delta t) = \max_{\delta \in \mathcal{D}_k^+} y_k(x^b, x^r, \delta, y_0) = \max_{\delta \in \mathcal{D}_k} y_k(x^b, x^r, \delta, y_0)$$

for all $t \in \mathcal{T}_k$, where the inequality holds because the integral is nondecreasing in t , and the equalities follow from the first part of the proof. In summary, we have shown that

$$\max_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, t) \leq \max \left\{ \max_{\delta \in \mathcal{D}_k} y_{k-1}(x^b, x^r, \delta, y_0), \max_{\delta \in \mathcal{D}_k} y_k(x^b, x^r, \delta, y_0) \right\}$$

for all $t \in \mathcal{T}_k$ and $k \in \mathcal{K}$. This implies that

$$\max_{t \in \mathcal{T}} \max_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, t) \leq \max_{k \in \mathcal{K} \cup \{0\}} \max_{\delta \in \mathcal{D}_k} y_k(x^b, x^r, \delta, y_0).$$

On the other hand, we have

$$\max_{t \in \mathcal{T}} \max_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, t) \geq \max_{k \in \mathcal{K} \cup \{0\}} \max_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, k\Delta t) = \max_{k \in \mathcal{K} \cup \{0\}} \max_{\delta \in \mathcal{D}_k} y_k(x^b, x^r, \delta, y_0),$$

where the equality follows from the first part of the proof. Combining the above inequalities implies (9), and thus the claim follows. \square

Proposition 8. *The following equivalence holds.*

$$y(x^b, x^r, \delta, y_0, t) \geq \underline{y} \quad \forall \delta \in \mathcal{D}, \forall t \in \mathcal{T} \iff y_k(x^b, x^r, \delta, y_0) \geq \underline{y} \quad \forall \delta \in \mathcal{D}_k, \forall k \in \mathcal{K} \cup \{0\}$$

The proof of Proposition 8 is significantly more challenging than that of Proposition 7 because $y(x^b, x^r, \delta, y_0, t)$ is concave in δ . We make it more digestible by first proving two lemmas.

Lemma 1. If $f : \mathbb{R} \times \mathcal{T} \rightarrow \mathbb{R}$ is concave, continuous and nonincreasing in its first argument and piecewise constant on the trading intervals \mathcal{T}_k , $k \in \mathcal{K}$, in its second argument, then

$$\min_{\delta \in \mathcal{D}^+} \int_0^t f(\delta(t'), t') dt' = \min_{\delta \in \mathcal{D}^+ \cap \mathcal{L}(\mathcal{T}, \{0,1\})} \int_0^t f(\delta(t'), t') dt' \quad \forall t \in \mathcal{T}.$$

Proof. For ease of exposition, assume first that $t = T$ and define $t_k = \Delta t(k-1)$ for every $k \in \mathcal{K}$. For every approximation parameter $N \in \mathbb{N}$ we define $\mathcal{N} = \{1, \dots, N\}$ and set

$$\mathcal{T}_{k,n}^N = \left[\Delta t \left(k-1 + \frac{n-1}{N} \right), \Delta t \left(k-1 + \frac{n}{N} \right) \right) \quad \forall k \in \mathcal{K}, \forall n \in \mathcal{N}.$$

Note that the $\mathcal{T}_{k,n}^N$, $n \in \mathcal{N}$, are mutually disjoint and that their union coincides with the k -th trading interval \mathcal{T}_k . Next, introduce a lifting operator $L_N : \mathbb{R}^{K \times N} \rightarrow \mathcal{L}(\mathcal{T}, \mathbb{R})$ defined through $(L_N \delta)(t) = \delta_{k,n}$ if $t \in \mathcal{T}_{k,n}^N$ for $k \in \mathcal{K}$ and $n \in \mathcal{N}$. In addition, let $L_N^\dagger : \mathcal{L}(\mathcal{T}, \mathbb{R}) \rightarrow \mathbb{R}^{K \times N}$ be the

corresponding adjoint operator defined through $(L_N^\dagger \delta)_{k,n} = \frac{N}{\Delta t} \int_{\mathcal{T}_{k,n}^N} \delta(t) dt$ for $k \in \mathcal{K}$ and $n \in \mathcal{N}$.

Using this notation, we first prove that

$$\lim_{N \rightarrow \infty} \sum_{k \in \mathcal{K}} \sum_{n \in \mathcal{N}} f \left((L_N^\dagger \delta)_{k,n}, t_k \right) \frac{\Delta t}{N} = \int_0^T f(\delta(t), t) dt \quad (11)$$

for any fixed $\delta \in \mathcal{D}^+$. As f is continuous and nonincreasing in its first argument and piecewise constant in its second argument, we have

$$\inf_{t \in \mathcal{T}_{k,n}^N} f(\delta(t), t) = f \left(\sup_{t \in \mathcal{T}_{k,n}^N} \delta(t), t_k \right) \leq f \left((L_N^\dagger \delta)_{k,n}, t_k \right) \leq f \left(\inf_{t \in \mathcal{T}_{k,n}^N} \delta(t), t_k \right) = \sup_{t \in \mathcal{T}_{k,n}^N} f(\delta(t), t)$$

for every $k \in \mathcal{K}$, $n \in \mathcal{N}$ and $N \in \mathbb{N}$. Summing over k and n thus yields

$$\sum_{k \in \mathcal{K}} \sum_{n \in \mathcal{N}} \inf_{t \in \mathcal{T}_{k,n}^N} f(\delta(t), t) \frac{\Delta t}{N} \leq \sum_{k \in \mathcal{K}} \sum_{n \in \mathcal{N}} f \left((L_N^\dagger \delta)_{k,n}, t_k \right) \frac{\Delta t}{N} \leq \sum_{k \in \mathcal{K}} \sum_{n \in \mathcal{N}} \sup_{t \in \mathcal{T}_{k,n}^N} f(\delta(t), t) \frac{\Delta t}{N}$$

for every $N \in \mathbb{N}$. As $f(\delta(t), t)$ constitutes a composition of a continuous function with a Riemann integrable function, it is also Riemann integrable. Thus, the lower and upper Riemann sums in the above inequality both converge to $\int_0^T f(\delta(t), t) dt$ as N tends to infinity. This observation establishes (11). As $\delta \in \mathcal{D}^+$ was chosen arbitrarily, we may thus conclude that

$$\inf_{\delta \in \mathcal{D}^+} \int_0^T f(\delta(t), t) dt = \inf_{\delta \in \mathcal{D}^+} \lim_{N \rightarrow \infty} \sum_{k \in \mathcal{K}} \sum_{n \in \mathcal{N}} f \left((L_N^\dagger \delta)_{k,n}, t_k \right) \frac{\Delta t}{N}.$$

For the following derivations we introduce the auxiliary uncertainty set

$$\mathcal{D}_{KN}^+ = \left\{ \boldsymbol{\delta} \in [-1, 1]^{KN} : \sum_{l=1+\lceil m-N\Gamma/\Delta t \rceil}^m \delta_l \leq N \frac{\gamma}{\Delta t} \quad \forall m = 1, \dots, KN \right\}$$

for $N \in \mathbb{N}$. By slight abuse of notation, we henceforth naturally identify any matrix $\boldsymbol{\delta} \in \mathbb{R}^{K \times N}$ with the vector obtained by concatenating the rows of $\boldsymbol{\delta}$. This convention allows us, for example, to write $\boldsymbol{\delta} \in \mathcal{D}_{KN}^+$ even if $\boldsymbol{\delta}$ was initially defined as a $K \times N$ -matrix. By repeating the arguments of Proposition 3, it is easy to show that $L_N \mathcal{D}_{KN}^+ \subseteq \mathcal{D}^+$ and $L_N^\dagger \mathcal{D}^+ = \mathcal{D}_{KN}^+$ for all $N \in \mathbb{N}$. Using these relations, we will now prove that

$$\inf_{\delta \in \mathcal{D}^+} \int_0^T f(\delta(t), t) dt = \lim_{N \rightarrow \infty} \inf_{\boldsymbol{\delta} \in \mathcal{D}_{KN}^+} \sum_{k \in \mathcal{K}} \sum_{n \in \mathcal{N}} f(\delta_{k,n}, t_k) \frac{\Delta t}{N}. \quad (12)$$

To this end, select any $\epsilon > 0$ and $\delta^* \in \mathcal{D}^+$ with $\int_0^T f(\delta^*(t), t) dt \leq \inf_{\delta \in \mathcal{D}^+} \int_0^T f(\delta(t), t) dt + \epsilon$, and choose N_ϵ large enough such that

$$\left| \sum_{k \in \mathcal{K}} \sum_{n \in \mathcal{N}} f \left((L_N^\dagger \delta^*)_{k,n}, t_k \right) \frac{\Delta t}{N} - \int_0^T f(\delta^*(t), t) dt \right| \leq \epsilon \quad \forall N \geq N_\epsilon.$$

Note that such an N_ϵ exists thanks to (11). For any $N \geq N_\epsilon$, we thus find

$$\begin{aligned}
0 &\leq \inf_{\delta \in \mathcal{D}_{KN}^+} \sum_{k \in \mathcal{K}} \sum_{n \in \mathcal{N}} f(\delta_{k,n}, t_k) \frac{\Delta t}{N} - \inf_{\delta \in \mathcal{D}^+} \int_0^T f(\delta(t), t) dt \\
&\leq \inf_{\delta \in \mathcal{D}_{KN}^+} \sum_{k \in \mathcal{K}} \sum_{n \in \mathcal{N}} f(\delta_{k,n}, t_k) \frac{\Delta t}{N} - \int_0^T f(\delta^*(t), t) dt + \epsilon \\
&\leq \sum_{k \in \mathcal{K}} \sum_{n \in \mathcal{N}} f\left(\left(L_N^\dagger \delta^*\right)_{k,n}, t_k\right) \frac{\Delta t}{N} - \int_0^T f(\delta^*(t), t) dt + \epsilon \leq 2\epsilon,
\end{aligned}$$

where the first inequality holds because $L_N \mathcal{D}_{KN}^+ \subseteq \mathcal{D}^+$, the second inequality follows from the choice of δ^* , the third inequality exploits the identity $L_N^\dagger \mathcal{D}^+ = \mathcal{D}_{KN}^+$, and the fourth inequality holds because $N \geq N_\epsilon$. As $\epsilon > 0$ was chosen arbitrarily, Equation (12) follows.

In order to prove that

$$\inf_{\delta \in \mathcal{D}^+} \int_0^T f(\delta(t), t) dt = \inf_{\delta \in \mathcal{D}^+ \cap \mathcal{L}(\mathcal{T}, \{0,1\})} \int_0^T f(\delta(t), t) dt, \quad (13)$$

we first observe that

$$\begin{aligned}
\inf_{\delta \in \mathcal{D}^+} \int_0^T f(\delta(t), t) dt &= \lim_{N \rightarrow \infty} \inf_{\delta \in \mathcal{D}_{KN}^+} \sum_{k \in \mathcal{K}} \sum_{n \in \mathcal{N}} f(\delta_{k,n}, t_k) \frac{\Delta t}{N} \\
&= \lim_{N \rightarrow \infty} \inf_{\delta \in \mathcal{D}_{KN}^+ \cap \{0,1\}^{KN}} \sum_{k \in \mathcal{K}} \sum_{n \in \mathcal{N}} f(\delta_{k,n}, t_k) \frac{\Delta t}{N}.
\end{aligned} \quad (14)$$

Here, the first equality follows from (12), and the second equality holds because f is concave in its first argument, which implies that the minimum over δ is attained at a vertex of the polyhedron \mathcal{D}_{KN}^+ . As all vertices of \mathcal{D}_{KN}^+ are binary by virtue of Lemma 2 below, we can restrict δ to $\{0,1\}^{K \times N}$ without loss of optimality. To prove (13), select any $\epsilon > 0$ and $N \in \mathbb{N}$ large enough such that

$$\left| \min_{\delta \in \mathcal{D}_{KN}^+ \cap \{0,1\}^{K \times N}} \sum_{k \in \mathcal{K}} \sum_{n \in \mathcal{N}} f(\delta_{k,n}, t_k) \frac{\Delta t}{N} - \inf_{\delta \in \mathcal{D}^+} \int_0^T f(\delta(t), t) dt \right| \leq \epsilon. \quad (15)$$

Note that such an N exists because of (14). Next, let δ^* be a minimizer of the discrete optimization problem on the left hand side of the above expression, and set $\delta^* = L_N \delta^*$. By (15) and because δ^* is constant on the intervals $\mathcal{T}_{k,n}^N$, we thus have

$$\left| \int_0^T f(\delta^*(t), t) dt - \inf_{\delta \in \mathcal{D}^+} \int_0^T f(\delta(t), t) dt \right| \leq \epsilon.$$

As $L_N \mathcal{D}_{KN}^+ \subseteq \mathcal{D}^+$ and $\delta^* \in \{0,1\}^{K \times N}$, we further have $\delta^* \in \mathcal{D}^+ \cap \mathcal{L}(\mathcal{T}, \{0,1\})$. As ϵ was chosen arbitrarily, Equation (13) follows.

If $t \in \mathcal{T}$ is a multiple of $1/(KN)$ for some $N \in \mathbb{N}$, then $f(\delta(t'), t')$ can be set to 0 for all $t' \geq t$, and the above proof remains valid with obvious minor modifications. For any other $t \in \mathcal{T}$, the claim follows from a continuity argument. Details are omitted for brevity. \square

Lemma 2. For any $N \in \mathbb{N}$, all vertices of the polyhedron

$$\mathcal{D}_{KN}^+ = \left\{ \boldsymbol{\delta} \in [0, 1]^{KN} : \sum_{l=1+[m-N\Gamma/\Delta t]^+}^m \delta_l \leq N \frac{\gamma}{\Delta t} \forall m = 1, \dots, KN \right\}$$

are binary vectors.

Proof. The polyhedron \mathcal{D}_{KN}^+ can be represented more concisely as $\{\boldsymbol{\delta} \in \mathbb{R}_+^{KN} : \mathbf{A}\boldsymbol{\delta} \leq \mathbf{b}\}$, where

$$\mathbf{A} = \begin{pmatrix} \mathbf{C} \\ \mathbf{I} \end{pmatrix} \in \mathbb{R}^{2KN \times KN}, \quad \mathbf{b} = \begin{bmatrix} N \frac{\gamma}{\Delta t} \mathbf{1} \\ \mathbf{1} \end{bmatrix} \in \mathbb{R}^{2KN}$$

and $\mathbf{C} \in \mathbb{R}^{KN \times KN}$ is defined through $C_{ij} = 1$ if $i - N\Gamma/\Delta t < j \leq i$ and $C_{ij} = 0$ otherwise. Here, \mathbf{I} denotes the identity matrix and $\mathbf{1}$ the column vector of 1s in \mathbb{R}^{KN} . By construction, \mathbf{A} is a binary matrix where the 1s appear consecutively in each row. Proposition 2.1 and Corollary 2.10 by Nemhauser and Wolsey (1999) thus imply that \mathbf{A} is totally unimodular. As $\mathbf{b} \in \mathbb{Z}^{KN}$ because of Assumption 2, all vertices of \mathcal{D}_{KN}^+ are integral thanks to Proposition 2.2 again by Nemhauser and Wolsey (1999). In addition, as $\mathcal{D}_{KN}^+ \subseteq [0, 1]^{KN}$, the vertices of \mathcal{D}_{KN}^+ are in fact binary vectors. \square

We are now ready to prove Proposition 8.

Proof of Proposition 8. The claim follows if we can show that

$$\min_{t \in \mathcal{T}} \min_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, t) = \min_{k \in \mathcal{K} \cup \{0\}} \min_{\delta \in \mathcal{D}_k} y_k(x^b, x^r, \delta, y_0). \quad (16)$$

In the first part of the proof, we reformulate the continuous non-convex minimization problem $\min_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, t)$ as a continuous linear program. To ease notation, we set $\Delta\eta = \frac{1}{\eta^-} - \eta^+ \geq 0$ and define the auxiliary functions

$$\chi(\delta(t), t) = \max \left\{ \eta^+ x^r(t) \delta(t), \frac{1}{\eta^-} x^r(t) \delta(t) - \Delta\eta x^b(t) \right\}$$

and

$$m(t) = \max \left\{ \eta^+ x^r(t), \frac{1}{\eta^-} x^r(t) - \Delta\eta x^b(t) \right\}$$

for all $t \in \mathcal{T}$. The function $\chi(\delta(t), t)$ can be viewed as a *nonlinear decision rule* of the uncertain frequency deviation $\delta(t)$. Using these conventions, we find

$$\begin{aligned}
\min_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, t) &= \min_{\delta \in \mathcal{D}^+} y(x^b, x^r, -\delta, y_0, t) \\
&= y_0 + \min_{\delta \in \mathcal{D}^+} \int_0^t \eta^+ x^b(t') - \chi(\delta(t'), t') - d(t') dt' \\
&= y_0 + \min_{\delta \in \mathcal{D}^+ \cap \mathcal{L}(\mathcal{T}, \{0,1\})} \int_0^t \eta^+ x^b(t') - \chi(\delta(t'), t') - d(t') dt' \quad (17) \\
&= y_0 + \min_{\delta \in \mathcal{D}^+ \cap \mathcal{L}(\mathcal{T}, \{0,1\})} \int_0^t \eta^+ x^b(t') - m(t')\delta(t') - d(t') dt' \\
&= y_0 + \min_{\delta \in \mathcal{D}^+} \int_0^t \eta^+ x^b(t') - m(t')\delta(t') - d(t') dt',
\end{aligned}$$

where the first equality holds because the statements $\delta \in \mathcal{D}$, $-\delta \in \mathcal{D}$ and $|\delta| \in \mathcal{D}^+$ are all equivalent and because y is nondecreasing in δ thanks to Proposition 1. The second equality follows from the definitions of y , y^+ , y^- and χ , and the third equality is a direct consequence of Lemma 1, which applies because $-\chi$ is concave and nonincreasing in its first argument and, by virtue of Assumption 1, piecewise constant in its second argument. The fourth equality holds because $\chi(\delta(t'), t') = m(t')\delta(t')$ whenever $\delta(t') \in \{0, 1\}$, and the last equality follows again from Lemma 1. Note that $m(t')\delta(t')$ constitutes a *linear decision rule* of $\delta(t)$.

In the second part of the proof we assume that $t = k\Delta t$ for some $k \in \mathcal{K} \cup \{0\}$ and show that

$$\min_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, t) = \min_{\delta \in \mathcal{D}_{\mathcal{K}}} y_k(\mathbf{x}^b, \mathbf{x}^r, \boldsymbol{\delta}, y_0).$$

To this end, we define $\chi_l(\delta_l) = \max\{\eta^+ x_l^r \delta_l, \frac{1}{\eta^-} x_l^r \delta_l - \Delta \eta x_l^b\}$ and $m_l = \max\{\eta^+ x_l^r, \frac{x_l^r}{\eta^-} - \Delta \eta x_l^b\}$ for all $l \in \mathcal{K}$. By (17), we thus have

$$\begin{aligned}
\min_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, k\Delta t) &= y_0 + \min_{\delta \in \mathcal{D}^+} \int_0^{k\Delta t} \eta^+ x^b(t') - m(t')\delta(t') - d(t') dt' \\
&= y_0 + \min_{\delta \in \mathcal{D}^+} \sum_{l=1}^k \int_{\mathcal{T}_l} \eta^+ x_l^b - m_l \delta(t') - d_l dt' \quad (18) \\
&= y_0 + \min_{\delta \in \mathcal{D}_{\mathcal{K}}^+} \Delta t \sum_{l=1}^k \eta^+ x_l^b - m_l \delta_l - d_l \\
&= \min_{\delta \in \mathcal{D}_{\mathcal{K}}} y_k(\mathbf{x}^b, \mathbf{x}^r, \boldsymbol{\delta}, y_0),
\end{aligned}$$

where the second equality holds because d , x^b and x^r are piecewise constant by virtue of Assumption 1, which implies that $m(t') = m_l$ for every $t' \in \mathcal{T}_l$. The third equality then follows from

Proposition 3. The fourth equality can be proved by reversing the arguments from (17) with obvious minor modifications. In fact, as the frequency deviation scenarios are now piecewise constant and can be encoded by finite-dimensional vectors, the proof requires no cumbersome limiting arguments as the ones developed in the proof of Lemma 1. We omit the details for brevity.

In the third part of the proof we assume that $t \in \mathcal{T}_k$ for some $k \in \mathcal{K}$ and show that

$$\min_{t \in \mathcal{T}_k} \min_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, t) = \min_{l \in \{k-1, k\}} \min_{\delta \in \mathcal{D}_{\mathcal{K}}} y_l(\mathbf{x}^b, \mathbf{x}^r, \boldsymbol{\delta}, y_0).$$

As in the proof of Proposition 7, we distinguish whether or not the vehicle is driving in period \mathcal{T}_k . Specifically, if the vehicle is driving in period \mathcal{T}_k , then $\bar{y}^+(t) = \bar{y}^-(t) = 0$, which implies that

$$\begin{aligned} \min_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, t) &= \min_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, (k-1)\Delta t) - \int_{(k-1)\Delta t}^t d(t') dt \\ &= \min_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, k\Delta t) + \int_t^{k\Delta t} d(t') dt \\ &\geq \min_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, k\Delta t) = \min_{\delta \in \mathcal{D}_{\mathcal{K}}} y_k(\mathbf{x}^b, \mathbf{x}^r, \boldsymbol{\delta}, y_0). \end{aligned}$$

Here, the inequality holds because $d(t) \geq 0$ for all $t \in \mathcal{T}_k$, and the last equality follows from (18).

Otherwise, if the vehicle is parked in period \mathcal{T}_k , then $d(t) = 0$ for all $t \in \mathcal{T}_k$, and hence

$$\begin{aligned} &\min_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, t) \\ &= y_0 + \min_{\delta \in \mathcal{D}^+} \sum_{l=1}^{k-1} \int_{\mathcal{T}_l} \eta^+ x_l^b - m_l \delta(t') - d_l dt' + \int_{(k-1)\Delta t}^t \eta^+ x_k^b - m_k \delta(t') dt' \\ &= y_0 + \min_{\delta \in \mathcal{D}^+(t)} \Delta t \sum_{l=1}^{k-1} \eta^+ x_l^b - d_l + (t - (k-1)\Delta t) \eta^+ x_k^b - \sum_{l=1}^k \int_{\mathcal{T}_l} m_l \delta(t') dt' \\ &= y_0 + \Delta t \sum_{l=1}^{k-1} \eta^+ x_l^b - d_l + (t - (k-1)\Delta t) \eta^+ x_k^b - \max_{\delta \in \mathcal{D}_{\mathcal{K}}^+(t)} \Delta t \sum_{l=1}^k m_l \delta_l, \end{aligned} \tag{19}$$

where we use the time-dependent uncertainty sets

$$\mathcal{D}^+(t) = \left\{ \delta \in \mathcal{D}^+ : \delta(t') = 0 \forall t' \in [t, k\Delta t] \right\} \quad \text{and} \quad \mathcal{D}_{\mathcal{K}}^+(t) = \left\{ \delta \in \mathcal{D}_{\mathcal{K}}^+ : \delta_k \leq \frac{t - (k-1)\Delta t}{\Delta t} \right\}$$

to simplify the notation. The first equality in (19) follows from (17) and Assumption 1. Note that $\delta(t')$ does not impact the objective function of the resulting minimization problem over \mathcal{D}^+ for any $t' > t$. It is therefore optimal to set $\delta(t') = 0$ for all $t' \geq t$ and, in particular, for all $t' \in [t, k\Delta t]$. This restriction has no impact on the objective function but maximizes nature's

flexibility in selecting harmful frequency deviations $\delta(t')$ for $t' \leq t$. Hence, the second equality in (19) follows. As δ is now integrated against a piecewise constant function, it may be averaged over the trading intervals without changing its objective function value. The third equality in (19) thus holds because $L^\dagger \mathcal{D}^+(t) = \mathcal{D}_K^+(t)$, which can be proved similarly to Proposition 3 by noting that

$$(\mathcal{L}^\dagger \delta)_k = \frac{1}{\Delta t} \int_{(k-1)\Delta t}^t \delta(t') dt' \leq \frac{t - (k-1)\Delta t}{\Delta t} \quad \forall \delta \in \mathcal{D}^+(t).$$

In the following we show that (19) is piecewise affine in t . To this end, note that the optimization problem in the last line of (19) can be expressed more concisely as the standard form linear program

$$\begin{aligned} \min_{z \geq \mathbf{0}} \quad & \mathbf{c}^\top z \\ \text{s.t.} \quad & \mathbf{A}z = \mathbf{b}(t), \end{aligned} \tag{20}$$

where $\mathbf{z}^\top = (\boldsymbol{\delta}^\top, \mathbf{s}^\top) \in \mathbb{R}^K \times \mathbb{R}^{2K}$ combines the (averaged) frequency deviations in the trading intervals with a vector of slack variables. Here, the vector $\mathbf{c} \in \mathbb{R}^{3K}$ of objective function coefficients is defined through $c_l = -m_l \Delta t$ if $l \leq k$ and $c_l = 0$ otherwise. The constraints involve the matrix

$$\mathbf{A} = \begin{pmatrix} \mathbf{C} & \mathbf{I} & \mathbf{0} \\ \mathbf{I} & \mathbf{0} & \mathbf{I} \end{pmatrix} \in \mathbb{R}^{2K \times 3K},$$

where $\mathbf{C} \in \mathbb{R}^{K \times K}$ is defined through $C_{ij} = 1$ if $i - \Gamma/\Delta t < j \leq i$ and $C_{ij} = 0$ otherwise, and the vector $\mathbf{b}(t) \in \mathbb{R}^{2K}$ is defined through $b_l(t) = \frac{t - (k-1)\Delta t}{\Delta t}$ if $l = k + K$ and $b_l = 1$ otherwise. By Lemma 2 and Proposition 2.1 of Nemhauser and Wolsey (1999), \mathbf{A} is totally unimodular.

Note that (20) is solvable for every $t \in \mathcal{T}_k$ because its feasible set is non-empty and compact. Next, choose any t_0 in the interior of \mathcal{T}_k , denote by \mathbf{B} an optimal basis matrix for problem (20) at $t = t_0$, and define $\mathbf{z}^*(t) = \mathbf{B}^{-1} \mathbf{b}(t)$ for all $t \in \mathcal{T}_k$. In the following, we will use local sensitivity analysis of linear programming to show that $\mathbf{z}^*(t)$ is optimal in (20) for all $t \in \mathcal{T}_k$. As the basis \mathbf{B} remains dual feasible when t deviates from t_0 , it suffices to show that

$$\mathbf{z}^*(t) = \mathbf{z}^*(t_0) + \frac{t - t_0}{\Delta t} \mathbf{B}^{-1} \mathbf{e}_{K+k} \geq \mathbf{0} \quad \forall t \in \mathcal{T}_k, \tag{21}$$

where \mathbf{e}_{K+k} denotes the $(K+k)$ -th standard basis vector in \mathbb{R}^{2K} (Bertsimas and Tsitsiklis, 1997, p. 207). To this end, note that \mathbf{B} is a non-singular square matrix constructed from $2K$ columns of \mathbf{A} and is therefore also totally unimodular. Moreover, \mathbf{B}^{-1} is totally unimodular because pivot operations preserve total unimodularity (Nemhauser and Wolsey, 1999, Proposition 2.1). Hence,

we have $\mathbf{B}^{-1}\mathbf{e}_{K+k} \in \{-1, 0, 1\}^{2K}$. By construction, we further have $\mathbf{b}(t) \in \{0, 1\}^{2K}$ for $t = k\Delta t$, which implies that $\mathbf{z}^*(k\Delta t) \in \mathbb{Z}^{2K}$. Evaluating (21) at $t = k\Delta t$ then yields

$$\mathbf{z}^*(t_0) = \mathbf{z}^*(k\Delta t) - \frac{k\Delta t - t_0}{\Delta t} \mathbf{B}^{-1} \mathbf{e}_{K+k},$$

which ensures that $\mathbf{z}^*(k\Delta t) \geq \mathbf{0}$. Indeed, if any component of the integral vector $\mathbf{z}^*(k\Delta t)$ was strictly negative, it would have to be smaller or equal to -1 . As t_0 resides in the interior of \mathcal{T}_k and thus $|(k\Delta t - t_0)/\Delta t| < 1$, the corresponding component of $\mathbf{z}^*(t_0)$ would then also have to be strictly negative. This, however, contradicts the optimality of $\mathbf{z}^*(t_0)$, which implies that $\mathbf{z}^*(t_0) \geq \mathbf{0}$. Hence, we have $\mathbf{z}^*(k\Delta t) \geq \mathbf{0}$. One can use similar arguments to prove that $\mathbf{z}^*((k-1)\Delta t) \geq \mathbf{0}$. As $\mathbf{z}^*(t)$ is affine in t , it is indeed non-negative for all $t \in \mathcal{T}_k$.

The above reasoning shows that $\mathbf{z}^*(t)$ is optimal in (20) and that the minimum of (20) is affine in t on \mathcal{T}_k . Equation (19) further implies that $\min_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, t)$ is affine in t on \mathcal{T}_k , and thus

$$\min_{t \in \mathcal{T}_k} \min_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, t) = \min_{l \in \{k-1, k\}} \min_{\delta \in \mathcal{D}} y(x^b, x^r, \delta, y_0, l\Delta t) = \min_{l \in \{k-1, k\}} \min_{\delta \in \mathcal{D}_K} y_l(\mathbf{x}^b, \mathbf{x}^r, \boldsymbol{\delta}, y_0),$$

where the second equality follows from (18). As $k \in \mathcal{K}$ was chosen arbitrarily, (16) follows. \square

Proposition 9. *The following equality holds.*

$$\max_{\delta \in \hat{\mathcal{D}}, y_0 \in \hat{\mathcal{Y}}_0} \varphi(y(x^b, x^r, \delta, y_0, T)) = \max_{\delta \in \hat{\mathcal{D}}_{\mathcal{K}}, y_0 \in \hat{\mathcal{Y}}_0} \varphi(y_K(\mathbf{x}^b, \mathbf{x}^r, \boldsymbol{\delta}, y_0))$$

Proof. By introducing an auxiliary epigraphical variable z , we find

$$\begin{aligned} \max_{\delta \in \hat{\mathcal{D}}, y_0 \in \hat{\mathcal{Y}}_0} \varphi(y(x^b, x^r, \delta, y_0, T)) &= \begin{cases} \min_z z \\ \text{s.t. } z \geq \max_{\delta \in \hat{\mathcal{D}}, y_0 \in \hat{\mathcal{Y}}_0} \varphi(y(x^b, x^r, \delta, y_0, T)) \end{cases} \\ &= \begin{cases} \min_z z \\ \text{s.t. } z \geq \max_{\delta \in \hat{\mathcal{D}}, y_0 \in \hat{\mathcal{Y}}_0} a_n y(x^b, x^r, \delta, y_0, T) + b_n \quad \forall n \in \mathcal{N} \end{cases} \quad (22) \\ &= \begin{cases} \min_z z \\ \text{s.t. } z \geq \max_{\delta \in \hat{\mathcal{D}}, y_0 \in \hat{\mathcal{Y}}_0} a_n y_K(\mathbf{x}^b, \mathbf{x}^r, \boldsymbol{\delta}, y_0) + b_n \quad \forall n \in \mathcal{N} \end{cases} \\ &= \max_{\delta \in \hat{\mathcal{D}}_{\mathcal{K}}, y_0 \in \hat{\mathcal{Y}}_0} \varphi(y_K(\mathbf{x}^b, \mathbf{x}^r, \boldsymbol{\delta}, y_0)). \end{aligned}$$

The second equality follows from the definition of φ and the third equality follows from Propositions 7 and 8, which apply since $\hat{\mathcal{D}}$ and $\hat{\mathcal{D}}_{\mathcal{K}}$ have the same structures as \mathcal{D} and $\mathcal{D}_{\mathcal{K}}$, respectively. \square

Proof of Theorem 1. The claim follows immediately from Propositions 6–9. \square

Proof of Theorem 2. By introducing embedded optimization problems that evaluate the (decision-dependent) worst-case frequency deviation scenarios and by replacing the uncertain initial state-of-charge in each robust constraint with its (decision-independent) worst-case value, $(\mathbf{R}_{\mathcal{K}})$ becomes

$$\begin{aligned}
& \min_{\mathbf{x}^b, \mathbf{x}^r \in \mathcal{X}_{\mathcal{K}}, z \in \mathbb{R}} c_{\mathcal{K}}(\mathbf{x}^b, \mathbf{x}^r) + z \\
& \text{s.t.} \quad \max_{\delta \in \mathcal{D}_{\mathcal{K}}} y^+(x_k^b, x_k^r, \delta_k) \leq \bar{y}_k^+ \quad \forall k \in \mathcal{K} & \text{(a)} \\
& \quad \max_{\delta \in \mathcal{D}_{\mathcal{K}}} y^-(x_k^b, x_k^r, \delta_k) \leq \bar{y}_k^- \quad \forall k \in \mathcal{K} & \text{(b)} \\
& \quad \max_{\delta \in \mathcal{D}_{\mathcal{K}}} y_k(\mathbf{x}^b, \mathbf{x}^r, \delta, \bar{y}_0) \leq \bar{y} \quad \forall k \in \mathcal{K} \cup \{0\} & \text{(c)} \\
& \quad \min_{\delta \in \mathcal{D}_{\mathcal{K}}} y_k(\mathbf{x}^b, \mathbf{x}^r, \delta, \underline{y}_0) \geq \underline{y} \quad \forall k \in \mathcal{K} \cup \{0\} & \text{(d)} \\
& \quad \max_{y_0 \in \mathcal{Y}_0} \max_{\delta \in \mathcal{D}_{\mathcal{K}}} \varphi(y_k(\mathbf{x}^b, \mathbf{x}^r, \delta, y_0)) \leq z & \text{(e)}
\end{aligned} \tag{23}$$

Here, the worst-case cost-to-go has been moved from the objective function to the constraints by introducing the auxiliary epigraphical variable z . To show that (23) is equivalent to $(\mathbf{R}'_{\mathcal{K}})$, we reuse several results derived for the proof of Theorem 1. First, by Equation (8) in the proof of Proposition 6 the maximum charging power in (23a) equals

$$\max_{\delta \in \mathcal{D}_{\mathcal{K}}} y^+(x_k^b, x_k^r, \delta_k) = x_k^r + x_k^b.$$

Using similar arguments, it can be shown that the maximum discharging power in (23b) reduces to

$$\max_{\delta \in \mathcal{D}_{\mathcal{K}}} y^-(x_k^b, x_k^r, \delta_k) = x_k^r - x_k^b.$$

Next, Equation (10) in the proof of Proposition 7 reveals that, for any $k \in \mathcal{K} \cup \{0\}$, the maximum state-of-charge in (23c) is given by

$$\max_{\delta \in \mathcal{D}_{\mathcal{K}}} y_k(\mathbf{x}^b, \mathbf{x}^r, \delta, \bar{y}_0) = \bar{y}_0 + \max_{\delta \in \mathcal{D}_{\mathcal{K}}^+} \Delta t \sum_{l=1}^k \eta^+(x_l^b + \delta_l x_l^r) - d_l.$$

Similarly, Equation (18) in the proof of Proposition 8 implies that, for every $k \in \mathcal{K} \cup \{0\}$, the minimum state-of-charge in (23d) amounts to

$$\min_{\delta \in \mathcal{D}_{\mathcal{K}}} y_k(\mathbf{x}^b, \mathbf{x}^r, \delta, \underline{y}_0) = \underline{y}_0 + \min_{\delta \in \mathcal{D}_{\mathcal{K}}^+} \Delta t \sum_{l=1}^k \eta^+ x_l^b - m_l \delta_l - d_l,$$

where $m_l = \max\{\eta^+ x_l^r, \frac{1}{\eta^-} x_l^r - \Delta\eta x_l^b\}$ constitutes an implicit function of the market decisions x_l^b and x_l^r . As (23d) imposes a *lower* bound on the minimum state-of-charge, m_l may be reinterpreted as an auxiliary epigraphical variable that satisfies $m_l \geq \eta^+ x_l^r$ and $m_l \geq \frac{1}{\eta^-} x_l^r - \Delta\eta x_l^b$.

Finally, the maximum cost-to-go in (23e) can be reformulated as

$$\begin{aligned} \max_{y_0 \in \hat{\mathcal{Y}}_0} \max_{\delta \in \hat{\mathcal{D}}_{\mathcal{K}}} \varphi(y_K(\mathbf{x}^b, \mathbf{x}^r, \delta, y_0)) &= \max_{y_0 \in \hat{\mathcal{Y}}_0} \max_{\delta \in \hat{\mathcal{D}}_{\mathcal{K}}} \max_{n \in \mathcal{N}} a_n y_K(\mathbf{x}^b, \mathbf{x}^r, \delta, y_0) + b_n \\ &= \max_{\delta \in \hat{\mathcal{D}}_{\mathcal{K}}} \max \left\{ \max_{n \in \mathcal{N}_+} b_n + a_n y_K(\mathbf{x}^b, \mathbf{x}^r, \delta, \hat{y}_0), \max_{n \in \mathcal{N}_-} b_n + a_n y_K(\mathbf{x}^b, \mathbf{x}^r, \delta, \underline{y}_0), \max_{n \in \mathcal{N}_0} b_n \right\}, \end{aligned}$$

where the first equality follows from the definition of the convex piecewise affine function φ . The second equality holds because the order of maximization is immaterial, because $a_n > 0$ for all $n \in \mathcal{N}_+$, $a_n < 0$ for all $n \in \mathcal{N}_-$ and $a_n = 0$ for all $n \in \mathcal{N}_0$ and because the state-of-charge $y_K(\mathbf{x}^b, \mathbf{x}^r, \delta, y_0)$ increases with y_0 . Requiring the last expression to be smaller than or equal to z is equivalent to

$$\max_{y_0 \in \hat{\mathcal{Y}}_0} \max_{\delta \in \hat{\mathcal{D}}_{\mathcal{K}}} \varphi(y_K(\mathbf{x}^b, \mathbf{x}^r, \delta, y_0)) \leq z \iff \begin{cases} \max_{\delta \in \hat{\mathcal{D}}_{\mathcal{K}}} y_K(\mathbf{x}^b, \mathbf{x}^r, \delta, \hat{y}_0) \leq (z - b_n)/a_n & \forall n \in \mathcal{N}_+ \\ \min_{\delta \in \hat{\mathcal{D}}_{\mathcal{K}}} y_K(\mathbf{x}^b, \mathbf{x}^r, \delta, \underline{y}_0) \geq (z - b_n)/a_n & \forall n \in \mathcal{N}_- \\ b_n \leq z & \forall n \in \mathcal{N}_0. \end{cases}$$

As $\mathcal{D}_{\mathcal{K}}$ and $\hat{\mathcal{D}}_{\mathcal{K}}$ have the same structure, the embedded optimization problems over $\delta \in \hat{\mathcal{D}}_{\mathcal{K}}$ admit similar linear reformulations as the embedded optimization problems over $\delta \in \mathcal{D}_{\mathcal{K}}$ in (23c) and (23d). Substituting all obtained reformulations into (23) yields $(\mathbf{R}'_{\mathcal{K}})$. \square

Proof of Theorem 3. Problem $(\mathbf{R}'_{\mathcal{K}})$ can be reformulated as a linear program by using the standard machinery of robust optimization (Bertsimas and Sim, 2004; Ben-Tal et al., 2009). For example, the robust upper bound on the state-of-charge for a fixed $k \in \mathcal{K} \cup \{0\}$ is equivalent to

$$\bar{y}_0 + \Delta t \sum_{l=1}^k \eta^+ (x_l^b + \delta_l x_l^r) - d_l \leq \bar{y} \quad \forall \delta \in \mathcal{D}_{\mathcal{K}}^+ \iff \max_{\delta \in \mathcal{D}_{\mathcal{K}}^+} \Delta t \sum_{l=1}^k \eta^+ (x_l^b + \delta_l x_l^r) - d_l \leq \bar{y} - \bar{y}_0. \quad (24)$$

By strong linear programming duality, the maximization problem in (24) is equivalent to

$$\begin{aligned} \min_{\Lambda^+, \Theta^+ \in \mathbb{R}_+^{K \times K}} \quad & \sum_{l=1}^k \Delta t \left(\eta^+ x_l^b + \Lambda_{k,l}^+ - d_l \right) + \gamma \Theta_{k,l}^+ \\ \text{s.t.} \quad & \Lambda_{k,l}^+ + \sum_{i=l}^{j(k,l)} \Theta_{k,i}^+ \geq \eta^+ x_l^r \quad \forall l \in \mathcal{K} : l \leq k, \end{aligned} \quad (25)$$

and the minimum of (25) is smaller or equal to $\bar{y} - \bar{y}_0$ if and only if problem (25) has a feasible solution whose objective value is smaller or equal to $\bar{y} - \bar{y}_0$. Therefore, the robust constraint (24) is equivalent to the following system of ordinary linear constraints.

$$\begin{aligned} \sum_{l=1}^k \Delta t \left(\eta^+ x_l^b + \Lambda_{k,l}^+ - d_l \right) + \gamma \Theta_{k,l}^+ &\leq \bar{y} - \bar{y}_0 \\ \Lambda_{k,l}^+ + \sum_{i=l}^{I(k,l)} \Theta_{k,i}^+ &\geq \eta^+ x_l^r \quad \forall l \in \mathcal{K} : l \leq k \end{aligned}$$

The remaining robust constraints in $(R'_{\mathcal{K}})$ can be simplified in a similar manner. \square

Proof of Proposition 5. We prove the claim by backward induction. Note first that φ_H is convex and piecewise affine by definition. Next, fix any $h \in \mathcal{H}$, and assume that φ_{h+1} is convex and piecewise affine, which implies that problem (5) is structurally equivalent to problem (R). By Theorems 1, 2 and 3, problem (5) can therefore be reformulated as a linear program whose right hand side coefficients depend affinely on y_h . Global sensitivity analysis of linear programming (Bertsimas and Tsitsiklis, 1997, p. 214) then ensures that $\varphi_h(y_h)$ is convex and piecewise affine in y_h . \square