



Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

An Exponential Cone Programming Approach for Managing Electric Vehicle Charging

Li Chen; , Long He, Yangfang (Helen) Zhou

To cite this article:

Li Chen; , Long He, Yangfang (Helen) Zhou (2024) An Exponential Cone Programming Approach for Managing Electric Vehicle Charging. *Operations Research* 72(5):2215–2240. <https://doi.org/10.1287/opre.2023.2460>

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. You are free to download this work and share with others, but cannot change in any way or use commercially without permission, and you must attribute this work as “*Operations Research*. Copyright © 2023 The Author(s). <https://doi.org/10.1287/opre.2023.2460>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by-nc-nd/4.0/>.”

Copyright © 2023 The Author(s)

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Methods

An Exponential Cone Programming Approach for Managing Electric Vehicle Charging

Li Chen,^{a,b,*} Long He,^c Yangfang (Helen) Zhou^d

^aInstitute of Operations Research and Analytics, National University of Singapore, Singapore 117602; ^bDiscipline of Business Analytics, The University of Sydney, Sydney, New South Wales 2006, Australia; ^cSchool of Business, The George Washington University, Washington, DC 20052; ^dLee Kong Chian School of Business, Singapore Management University, Singapore 178899

*Corresponding author

Contact: chen_l@u.nus.edu,  <https://orcid.org/0000-0002-5370-8518> (LC); longhe@gwu.edu,  <https://orcid.org/0000-0002-9951-7450> (LH); helenzhou@smu.edu.sg,  <https://orcid.org/0000-0001-6408-3448> (Y(H)Z)

Received: January 4, 2021

Revised: November 23, 2021; January 5, 2023

Accepted: March 16, 2023

Published Online in Articles in Advance: May 16, 2023, and updated December 27, 2024

Area of Review: Optimization

<https://doi.org/10.1287/opre.2023.2460>

Copyright: © 2023 The Author(s)

Abstract. To support the rapid growth in global electric vehicle adoption, public charging of electric vehicles is crucial. We study the problem of an electric vehicle charging service provider, which faces (1) stochastic arrival of customers with distinctive arrival/departure times and energy requirements and (2) a total electricity cost including demand charges, which are costs related to the highest per-period electricity used in a finite horizon. We formulate its problem of scheduling vehicle charging to minimize the expected total cost as a stochastic program (SP). As this SP is large-scale, we solve it using exponential cone program (ECP) approximations. For the SP with unlimited chargers, we derive an ECP as an upper bound and characterize the bound on the gap between their theoretical performances. For the SP with limited chargers, we then extend this ECP by also leveraging the idea from distributionally robust optimization (DRO) of using an entropic dominance ambiguity set. Instead of using DRO to mitigate distributional ambiguity, we use it to derive an ECP as a tractable upper bound of the SP. We benchmark our ECP approach with sample average approximation (SAA) and a DRO approach using a semidefinite program (SDP) on numerical instances calibrated to real data. As our numerical instances are large-scale, we find that although SDP cannot be solved, ECP scales well and runs efficiently (about 50 times faster than SAA) and consequently results in a lower mean total cost than SAA. We then show that our ECP continues to perform well considering practical implementation issues, including a data-driven setting and an adaptive charging environment. We finally extend our ECP approaches (for both the uncapacitated and capacitated cases) to include the pricing decision and propose an alternating optimization algorithm, which performs better than SAA on our numerical instances. Our method of constructing ECPs can be potentially applicable to approximate more general two-stage linear SPs with fixed recourse. We also use ECP to generate managerial insights for both charging service providers and policymakers.

 **Open Access Statement:** This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. You are free to download this work and share with others, but cannot change in any way or use commercially without permission, and you must attribute this work as “Operations Research. Copyright © 2023 The Author(s). <https://doi.org/10.1287/opre.2023.2460>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by-nc-nd/4.0/>.”

Funding: L. Chen gratefully acknowledges the financial support from the Singapore Ministry of Education with the 2019 Academic Research Fund Tier 3 [Grant MOE-2019-T3-1-010].

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/opre.2023.2460>.

Keywords: stochastic programming • exponential cone programming • electric vehicle • demand charge • robust optimization

1. Introduction

Electric vehicles (EVs) are considered among the most promising technologies to decarbonize the transportation sector, with market share expected to grow from 1% in 2018 to about 30% of vehicle sales worldwide by 2030 (J. P. Morgan 2018). A key to the mass adoption of EVs is the ease of charging, where public charging will play an increasingly important role, for example, in supporting

adopters without residential charging and reducing “range anxiety” (Wood et al. 2017, McKinsey & Company 2018). Examples of public charging service providers include EVgo, Tesla, and ChargePoint, which have the largest market shares in the United States. It is estimated that this charging market in the United States will grow to \$18.6 billion, and the number of charging points will expand to 40 million globally by 2030 (Bloomberg 2018).

A service provider of EV charging faces significant operational challenges. First, there are stochastic arrivals of customers who have distinctive arrival time, desired departure time, and charging requirements for their EVs, observations that we show using real data in the numerical study section. Second, the tariff structure for an EV charging service provider in the United States—a commercial electricity consumer—includes *demand charges*, that is, costs “based on the highest average electricity usage occurring within a defined time interval (usually 15 minutes) during a billing period” (National Renewable Energy Laboratory 2017). The total demand charge is the sum of that applicable to all periods during the entire horizon and that applicable to periods during only on-peak or mid-peak hours: Time-differentiated demand charges were developed to alleviate the issue that all-period demand charges may penalize users with peak load, which may not coincide with the system peak (Veall 1983, Hausman and Neufeld 1984). This total demand charge for an EV charging service provider can be as high as 70% of its total electricity cost (Chitkara et al. 2016).

In this paper, we study the problem of a service provider in scheduling EV charging with the goal of minimizing the total expected cost. We model it as a stochastic program (SP). We characterize the random number of arriving customers to follow Poisson distributions, commonly used for customer arrivals into a charging station (Mak et al. 2013, Schneider et al. 2017, He et al. 2021).

This SP is large-scale due to a large number of customer types which differ in arrival time, departure time, and charging quantity. Therefore, we propose to solve it approximately using a new approach called exponential cone programming (Chares 2009). It belongs to the class of conic programming (Boyd and Vandenberghe 2004) that generalizes linear programming to incorporate inequalities defined by an exponential cone—a three-dimensional convex cone involving exponentials and logarithms. Exponential cone programming is an attractive approach as it can be solved efficiently (in polynomial time) owing to the recent availability of solvers, such as MOSEK (MOSEK ApS 2020). There have been only a few attempts to use exponential cone programming in real applications, such as inventory management, manpower planning, and insurance reimbursement (See and Sim 2010, Zhu et al. 2021, Jaillet et al. 2022). Among these, only See and Sim (2010) involves solving an SP, and they use an approach based on exponential cone programming without explicitly formulating an exponential cone program (ECP). In addition, See and Sim (2010) requires the random variables in the SP to be sub-Gaussian, different from our paper, and thus cannot be directly applied to our EV charging problem to derive ECPs. Therefore, in this paper, we develop new methods to derive our ECPs to solve our SP approximately. In particular, we use moment-generating functions (MGFs) of the random variables to bound an order

statistic in the objective introduced by the demand charge.

We then benchmark our ECP approach to solve the SP with two common approaches, sample average approximation (SAA) and a distributionally robust optimization (DRO) approach using a semidefinite program (SDP). We compare these approaches on numerical instances calibrated to real data of electricity tariff and customer arrivals. We summarize our theoretical and numerical contributions as follows:

- We derive an ECP approximation for the case in which the provider has unlimited chargers labeled ECP-U (where U represents “uncapacitated”). We show that the optimal expected cost of ECP-U gives an upper bound of that of our uncapacitated SP. We also characterize a bound on the gap between the theoretical performances of ECP-U and the SP.

- We derive another ECP approximation for the case with finite chargers, labeled ECP-C (where C represents “capacitated”), by combining the method in deriving ECP-U and the idea from DRO of employing an entropic dominance ambiguity set (Chen et al. 2019). Instead of using DRO to mitigate distributional ambiguity, we use it to derive ECP-C as a tractable upper bound of the capacitated SP. We show that the optimal expected cost of ECP-C converges to a finite value bounded by that of ECP-U as the number of chargers tends to infinity.

- We extend our ECP approach to include the pricing decision (together with the charging scheduling decisions) for both the uncapacitated and capacitated cases. We show that our ECP approximations can be solved efficiently using an alternating optimization algorithm.

- As our numerical instances are large-scale (with about 80,000 random variables and 700,000 decision variables), we show that whereas DRO-based SDP cannot be solved, ECP-C runs efficiently (about 50 times faster than SAA), and results in not only a lower standard deviation of the total cost but also a lower mean of this cost than SAA because SAA cannot be solved to optimality within a reasonable amount of time. We also find the optimality gap of ECP-C to be at most around 4% on all instances. In addition, we show the superior performance of ECP over SAA and other benchmarks continues to hold considering three practical implementation issues: finer time discretization, a data-driven setting where there are estimation errors in arrival rates or model mis-specification of arrival distribution, and an adaptive EV charging environment. In the extended model with pricing decision, we show that our ECP approach continues to perform better than SAA.

- Using the ECP approach on our numerical instances also results in the following managerial insights. First, it is beneficial to incorporate customer departure information in scheduling optimization. Second, either an on-peak or mid-peak demand charge alone is not sufficient to smooth the electricity load, but an all-period

demand charge alone can smooth the electricity load as well as when all three types of demand charge exist.

Our paper has the following academic and practice relevance: First, although developed for the EV charging problem, our method of constructing ECPs can be potentially applicable to approximate more general two-stage linear SPs with fixed recourse. Second, our ECP-C approximation can be used by charging service providers in managing EV charging due to its superior performance. Third, our numerical results provide insights for charging service providers on the value of customer departure information and for regulatory bodies on designing tariff structures in smoothing electricity load.

The rest of the paper is organized as follows. After reviewing the literature in Section 2, we introduce our SP in Section 3 for scheduling EV charging. In Section 4, we first develop ECP-U for the case in which the charging service provider has unlimited chargers and provide its performance guarantee; we then develop ECP-C for the capacitated case. In Section 5, we consider two benchmark approaches to solving our SP model. In Section 6, we consider an extension of our model to jointly optimize the pricing and scheduling decisions and derive ECP approximations, which we solve by developing an alternating optimization algorithm. In Section 7, using numerical instances calibrated to real data, we compare the performance of ECP-C with other approaches and generate managerial insights using ECP-C. We conclude and discuss future work in Section 8. Any proof not in the main text can be found in the Online Appendix.

2. Literature Review

Our paper contributes to the literature on optimization with uncertainty. Common approaches include stochastic programming (Danzig 1955, Shapiro et al. 2009, Birge and Louveaux 2011), robust optimization (Goldfarb and Iyengar 2003; Bertsimas and Sim 2004; Ben-Tal et al. 2009, 2013), and, more recently, DRO (Delage and Ye 2010, Wiesemann et al. 2014, Esfahani and Kuhn 2018, Bertsimas et al. 2019, and the references therein). We build an SP model for the EV charging scheduling problem and specifically contribute to the literature of using the expected recourse function approximation in solving SPs (see Birge and Louveaux (2011) chapter 8, for example) by proposing a new approach based on ECPs. ECPs have been used in a limited number of papers, such as See and Sim (2010), Jaillet et al. (2022), and Zhu et al. (2021), but none of the existing methods to derive ECPs applies to our EV charging problem. In particular, we develop ECP-U by using an MGF to bound the largest order statistic introduced by the demand charge. This method can be potentially applied to approximate more general two-stage linear SP with fixed recourse. We then combine this method together with the idea in the DRO framework of using an entropic dominance

ambiguity set (Chen et al. 2019) to derive ECP-C. In addition, all earlier work solves ECPs approximately using second-order cones or cutting plane due to a lack of efficient solvers, whereas we solve ECPs exactly.

Our paper also contributes to the nascent literature in EV operations management, which focuses on problems such as the planning and operations of battery swapping stations (Mak et al. 2013, Schneider et al. 2017, Sun et al. 2019), the adoption of EVs (Avci et al. 2014, Lim et al. 2014), charging infrastructure planning (He et al. 2021), and vehicle-to-grid operations (Zhang et al. 2021). In particular, He et al. (2021) present an integrated model to jointly determine the size and location of EV charging stations for a fleet of EVs in a vehicle sharing context. Zhang et al. (2021) formulate a two-stage stochastic integer program to study service-zone and facility capacity planning and fleet management in EV sharing systems with vehicle-to-grid integration. Different from those papers, which primarily focus on infrastructure planning, we optimize for a charging schedule that caters to various customer types differentiated by arrival/departure times and charging requirements.

In the literature on EV charging management from the perspective of a charging service provider, Jin et al. (2013) and Zhang et al. (2014) consider the problem of scheduling EV charging in the presence of an energy storage system. Our paper differs from the first in that our EV arrivals are stochastic while theirs is deterministic and differs from the second in that we optimize the charging decisions while their charging rate is fixed. Jiang and Powell (2016) optimize the charging of one EV within a reservation window and consider the risk of the charging cost, whereas we model the setting where the provider charges multiple EVs with stochastic arrivals. Wu et al. (2022) study the problem of incentivizing customers to delay charging by designing a menu of price and departure time using a mechanism design approach, while we study the problem of optimizing the charging schedule with stochastic customer arrivals using a stochastic optimization approach. In addition, different from all these papers, we consider the setting where the provider buys electricity from a utility company with a tariff structure that includes demand charges.

Notation. We use boldface uppercase and lowercase characters to denote matrices and vectors, respectively. For example, $x \in \mathbb{R}^n$ means that x is an n -dimensional real vector. For a matrix $A \in \mathbb{R}^{m \times n}$, A' is the transpose. As usual, we use $\mathbf{1}$ to denote a vector of all ones and $\mathbf{0}$ to denote the vector or a matrix of all zeros. We define the inner product $X \bullet Y = \text{trace}(X'Y)$ for any $X, Y \in \mathbb{R}^{m \times n}$. By $x \geq y$, we mean that x is greater than y component-wise. We denote by $[N] \triangleq \{1, 2, \dots, N\}$ the set of positive running indices up to N . We use $\lfloor \cdot \rfloor$ ($\lceil \cdot \rceil$) as a floor (ceil) function that takes a real number as input and gives as output the greatest (smallest) integer less (greater) than

or equal to the input. We use $\mathcal{P}_0(\mathcal{Z})$ to represent the set of all probability distributions on the support set $\mathcal{Z}$. A random vector, $\tilde{z}$, is denoted with a tilde sign and we use $\tilde{z} \sim \mathbb{P}, \mathbb{P} \in \mathcal{P}_0(\mathbb{R}^{I_z})$ to define $\tilde{z}$ as an I_z -dimensional random variable with distribution $\mathbb{P}$. We use $\mathbb{E}_{\mathbb{P}}[\cdot]$ to signify the expectation with respect to $\mathbb{P}$.

3. A Stochastic Programming Model for Scheduling EV Charging

We model the operations of a charging service provider to minimize the total expected cost in a finite horizon of T periods. The service provider has a limited number of chargers, denoted by C . We assume customers do not enter a station that does not have chargers available. This assumption is because EVs (e.g., Tesla) are equipped with onboard or mobile navigation systems that can help route them only to stations with available chargers.

As the EV charging industry is evolving quickly, there are currently various business practices. We consider one common practice where the service provider requests from each customer both *desired departure time* and *quantity to charge* (i.e., charging requirement) to manage the charging more effectively. For example, the adaptive charging network framework developed at Caltech¹ and the corresponding commercialized charging management software developed by Powerflex (2021) allow drivers to provide their estimated departure time and energy request through a mobile app. Moreover, this setting echoes the literature in EV charging scheduling (Jin et al. 2013, Xu et al. 2016). Providing such departure time information is intuitively more valuable for the case of long-duration charging, such as a few hours in the workplace and residential charging, than with short-duration charging, such as 30 minutes in public charging stations at highway stops. However, our model is general to accommodate both short- and long-duration charging. Furthermore, we show using data in Section 7.5.1 even if the duration is short (such as 35 minutes on average), the value of providing departure information for the scheduling optimization is still significant, and this value increases rapidly when the duration increases. This alludes to the value of service providers requesting such departure information even when the charging duration is short.

EV Arrivals. To characterize the customer arrival profile in the aforementioned general business practice, we classify EV customers into V types according to the triple (s_v, τ_v, u_v) , where s_v is the arrival time, τ_v is the desired departure time ($s_v \leq \tau_v$), and u_v is the charging requirement for customer type $v \in [V]$. This definition of customer types is also consistent with real data used in Section 7, where customers are heterogeneous in all three dimensions.

Recent research on EV operations (Mak et al. 2013, Schneider et al. 2017, He et al. 2021) often characterize

customer arrivals at a charging (or battery swapping) station by Poisson distributions, which is empirically verified in He et al. (2021). Similarly, we assume the arrivals of customer types $v \in [V]$ with infinite chargers are independent Poisson random variables, and each has an arrival rate λ_v . However, as the charging capacity is finite, the arrivals of customer types v (denoted by $\tilde{z}_v$) are truncated Poisson random variables,² and thus they become dependent in this case. Let $\mathcal{V}_t \triangleq \{v \in [V] | s_v \leq t \leq \tau_v\}$ denote the set of customer types at the station in period t . Let $\tilde{z} \triangleq (\tilde{z}_v)_{v \in [V]}$ denote the vector of these truncated Poisson random variables drawn from the joint distribution denoted by $\mathbb{P}^C$ (where C represents “capacitated”) with arrival rate $\lambda \triangleq (\lambda_v)_{v \in [V]}$ and support $\mathcal{Z}$ defined as

$$\mathcal{Z} \triangleq \left\{ z \geq 0 \mid \sum_{v \in \mathcal{V}_t} z_v \leq C, \forall t \in [T] \right\}. \quad (1)$$

Decisions. The service provider decides on a menu-based schedule where the charging decision is adaptive to customer types. Specifically, the menu-based charging schedule is $x \triangleq (x_{v,t})_{v \in [V], t \in \mathcal{T}_v}$, where $x_{v,t}$ denotes in period $t \in \mathcal{T}_v$ the charging speed for customer type v (i.e., the quantity of electricity used to charge the EV of this type in this period) and $\mathcal{T}_v \triangleq \{s_v, \dots, \tau_v\}$ is the set of periods within the charging window of customer type v . Our modeling of the charging speed optimization is motivated by the observation in our data that the average charging speeds for most EVs vary and are *not* at their maximums, even for chargers with charging power greater than 22kW (see Section 7). The sequence of events in each period $t \in [T]$ is as follows: At the beginning of each period t , the charging service provider observes the realization of customer arrival $\tilde{z}_v$ with $s_v = t$, and charges $x_{v,t}$ to EVs of type $v \in \mathcal{V}_t$ by looking up the menu-based charging schedule x . At the end of period t , customers with desired departure time equal to t (i.e., $\tau_v = t$) finish their charging service and leave the charging station.³ The feasible set of x , denoted by $\mathcal{X}$, is given as follows:

$$\mathcal{X} \triangleq \left\{ x \mid \begin{array}{l} \sum_{t \in \mathcal{T}_v} \eta x_{v,t} = u_v \quad \forall v \in [V] \\ 0 \leq x_{v,t} \leq K/\eta \quad \forall v \in [V], t \in \mathcal{T}_v \end{array} \right\},$$

where the first constraint indicates that customer type v needs to fulfill the charging requirement u_v before τ_v and $\eta \in (0, 1]$ is the ratio of the quantity of electricity increased in the battery to the quantity of electricity used to charge the battery. The second constraint indicates that the charging speed is within the limit, where K (in energy unit/period) is the maximum amount of energy a battery can be charged in each period.⁴ For a well-defined type $v \in [V]$, it thus satisfies $u_v \leq K(\tau_v - s_v + 1)$, as otherwise the charging requirement cannot be met even with charging at the maximum speed for the entire stay duration.

Total Cost. The objective of the service provider is to minimize its expected total electricity cost, which includes demand charge as mentioned in Section 1. This demand charge depends on the maximum per-period electricity used within the horizon and is imposed on almost all commercial electricity users (such as charging stations) by their utility firms in the United States (Mills et al. 2008, National Renewable Energy Laboratory 2017). The rationale is for utility firms to recover the maintenance cost of standby equipment in the event of high load from users and also incentivize them to smooth the electricity load over time. Let d (in \$/(energy unit/period)) denote the unit demand charge; let $f_t(x, \tilde{z})$ denote the total electricity used to charge EVs of all customers in period t , that is,

$$f_t(x, \tilde{z}) = \sum_{v \in \mathcal{V}_t} x_{v,t} \tilde{z}_v.$$

Given $(x, \tilde{z})$, the total cost of charging all EVs within the horizon, denoted by $c(x, \tilde{z})$, is as follows:

$$c(x, \tilde{z}) \triangleq d \max_{t \in [T]} \{f_t(x, \tilde{z})\} + \sum_{s \in [T]} e_s f_s(x, \tilde{z}), \quad (2)$$

where the first term is the demand charge, and the second term is the energy charge, with e_t (in \$/energy unit) denoting the unit time-of-use (TOU) energy charge. “Demand charge” and “energy charge” are terminology used in the rate structure of utility firms, where “demand” differs from the usual meaning in “customer demand,” and “charge” differs from that used in “EV charging.” For simplicity of exposition, the demand charge in (2) consists of only an all-period demand charge, applicable to all periods in the finite horizon. We consider the case in which the total demand charge is the sum of multiple types of demand charge, such as demand charge in peak hours and mid-peak hours, in the numerical study in Section 7. Last, our model and theoretical results can be easily extended to the case where demand charge depends on the maximum electricity usage *per multiple periods* (see an example in Section 7.4.1).

Model Formulation. We formulate the problem of scheduling EV charging as an SP:

$$\min_{x \in \mathcal{X}} \mathbb{E}_{\mathbb{P}^C} [c(x, \tilde{z})], \quad (3)$$

where the expectation is over $\tilde{z}$. We denote an optimal solution to (3) by x^* and the optimal value of (3) by π^* , that is, $\pi^* = \mathbb{E}_{\mathbb{P}^C} [c(x^*, \tilde{z})]$.

Note that $c(x, \tilde{z})$ is a nonlinear function that can be linearized with auxiliary variables, so Problem (3) can be rewritten as a two-stage stochastic linear program with complete recourse. Two-stage SPs are in general #P-hard, as evaluating the expectation of the random objective function is #P-hard (Hanasusanto et al. 2016).

In addition, our SP Model (3) is large-scale due to its large number of random variables, V , and decision variables: V is the product of three terms, the number of possible arrival periods, departure periods (given an arrival period), and charging requirements; the latter is the product of V and the number of periods within the charging window of all types. As we show in our numerical study calibrated to real data, for a representative one-day planning horizon, the numbers of random variables and decision variables are about 80,000 and 700,000, respectively. We next propose an approach to solve our SP based on ECP approximation in Section 4, and benchmark ECP with two common approaches, SAA and DRO, which we discuss in Section 5.

4. ECP Approximations

In this section, we first consider the uncapacitated case by developing an ECP approximation and providing a theoretical performance guarantee. We then extend it to the capacitated case by using the idea from the framework of DRO.

As mentioned before, an ECP is a conic program with constraints that can be represented by exponential cones, which are three-dimensional convex cones involving exponentials and logarithms:

$$\begin{aligned} \mathcal{K}_{\text{exp}} \triangleq & \{(x_1, x_2, x_3) | x_1 \geq x_2 \exp(x_3/x_2), x_2 > 0\} \\ & \cup \{(x_1, 0, x_3) | x_1 \geq 0, x_3 \leq 0\}. \end{aligned}$$

ECP generalizes the widely used LP and second-order conic program (SOCP) (Chandrasekaran and Shah 2017). ECPs can be solved in polynomial time using an interior point algorithm with an efficient solver, such as MOSEK. Because the most difficult step in solving (3) is computing the expectation of the largest order statistic introduced by the demand charge (while the energy charge is linear in $\tilde{z}$), ECPs become a natural choice for the approximation of (3) because we use the MGF (involving exponential functions) of $\tilde{z}$ to bound this order statistic.

4.1. Uncapacitated Case

Let $\mathbb{P}^\infty$ denote the joint distribution of $\tilde{z}$ for $C = \infty$. Note that $\tilde{z}_v$ s for $v \in [V]$ are independent Poisson random variables. Given any $\theta \triangleq (\theta_v)_{v \in [V]}$, the MGF of $\tilde{z} \sim \mathbb{P}^\infty$ is

$$\begin{aligned} \mathbb{E}_{\mathbb{P}^\infty} \left[\exp \left(\sum_{v \in [V]} \theta_v \tilde{z}_v \right) \right] &= \prod_{v \in [V]} \mathbb{E}_{\mathbb{P}^\infty} [\exp(\theta_v \tilde{z}_v)] \\ &= \prod_{v \in [V]} \exp(\lambda_v (e^{\theta_v} - 1)), \end{aligned} \quad (4)$$

where the first equality is due to the independence of $\tilde{z}_v$'s and the second equality follows from the closed-form MGF expression of a Poisson random variable $\tilde{z} \sim$

$\mathbb{P}$ with arrival rate λ , that is, $\mathbb{E}_{\mathbb{P}}[e^{\theta \tilde{z}}] = e^{\lambda(e^{\theta} - 1)}$ for any θ . Because (4) involves exponential functions, we obtain an upper bound of π^* , the optimal value of (3), using an ECP as follows.

Proposition 1. *When $C = \infty$, the optimal value of ECP-U gives an upper bound of π^* :*

$$\inf_{x \in \mathcal{X}, \kappa, \gamma, \mu > 0, \xi, \zeta} d(\kappa + \gamma) + \sum_{s \in [T]} e_s f_s(\mathbf{x}, \boldsymbol{\lambda})$$

$$\text{s.t. } \sum_{v \in \mathcal{V}_t} x_{v,t} \lambda_v \leq \gamma \quad \forall t \in [T], \quad (5a)$$

$$\text{(ECP-U)} \quad \mu \exp(x_{v,t}/\mu) \leq \xi_{v,t} \quad \forall t \in [T], v \in \mathcal{V}_t, \quad (5b)$$

$$\mu \exp\left(\left(-\kappa + \sum_{v \in \mathcal{V}_t} \lambda_v (\xi_{v,t} - x_{v,t} - \mu)\right) / \mu\right) \leq \zeta_t \quad (5c)$$

$$\quad \forall t \in [T],$$

$$\sum_{t \in [T]} \zeta_t \leq \mu. \quad (5d)$$

Proof. To obtain an upper bound of π^* , we first obtain an upper bound of $\mathbb{E}_{\mathbb{P}^\infty}[\max_{t \in [T]} f_t(\mathbf{x}, \tilde{\mathbf{z}})]$:

$$\begin{aligned} & \mathbb{E}_{\mathbb{P}^\infty} \left[\max_{t \in [T]} f_t(\mathbf{x}, \tilde{\mathbf{z}}) \right] \\ &= \mathbb{E}_{\mathbb{P}^\infty} \left[\max_{t \in [T]} (f_t(\mathbf{x}, \tilde{\mathbf{z}}) - f_t(\mathbf{x}, \boldsymbol{\lambda}) + f_t(\mathbf{x}, \boldsymbol{\lambda})) \right] \\ &\leq \mathbb{E}_{\mathbb{P}^\infty} \left[\max_{t \in [T]} (f_t(\mathbf{x}, \tilde{\mathbf{z}}) - f_t(\mathbf{x}, \boldsymbol{\lambda})) + \max_{t \in [T]} f_t(\mathbf{x}, \boldsymbol{\lambda}) \right] \\ &= \mathbb{E}_{\mathbb{P}^\infty} \left[\max_{t \in [T]} (f_t(\mathbf{x}, \tilde{\mathbf{z}}) - f_t(\mathbf{x}, \boldsymbol{\lambda})) \right] + \max_{t \in [T]} f_t(\mathbf{x}, \boldsymbol{\lambda}). \end{aligned} \quad (6)$$

We then obtain an upper bound of the first term in (6) given any $\mu > 0$ as follows:

$$\begin{aligned} & \mathbb{E}_{\mathbb{P}^\infty} \left[\max_{t \in [T]} (f_t(\mathbf{x}, \tilde{\mathbf{z}}) - f_t(\mathbf{x}, \boldsymbol{\lambda})) \right] \\ &\leq \mu \ln \mathbb{E}_{\mathbb{P}^\infty} \left[\exp \left(\max_{t \in [T]} (f_t(\mathbf{x}, \tilde{\mathbf{z}}) - f_t(\mathbf{x}, \boldsymbol{\lambda})) / \mu \right) \right] \\ &\leq \mu \ln \mathbb{E}_{\mathbb{P}^\infty} \left[\sum_{t \in [T]} \exp((f_t(\mathbf{x}, \tilde{\mathbf{z}}) - f_t(\mathbf{x}, \boldsymbol{\lambda})) / \mu) \right] \\ &= \mu \ln \sum_{t \in [T]} \mathbb{E}_{\mathbb{P}^\infty} \left[\exp((f_t(\mathbf{x}, \tilde{\mathbf{z}}) - f_t(\mathbf{x}, \boldsymbol{\lambda})) / \mu) \right] \\ &= \mu \ln \sum_{t \in [T]} \mathbb{E}_{\mathbb{P}^\infty} \left[\exp \left(\sum_{v \in \mathcal{V}_t} \frac{x_{v,t}}{\mu} (\tilde{z}_v - \lambda_v) \right) \right] \\ &= \mu \ln \sum_{t \in [T]} \prod_{v \in \mathcal{V}_t} \mathbb{E}_{\mathbb{P}^\infty} \left[\exp \left(\frac{x_{v,t}}{\mu} \tilde{z}_v \right) \right] \exp \left(-\frac{x_{v,t}}{\mu} \lambda_v \right) \\ &= \mu \ln \sum_{t \in [T]} \exp \left(\sum_{v \in \mathcal{V}_t} \lambda_v (e^{x_{v,t}/\mu} - 1 - x_{v,t}/\mu) \right), \end{aligned} \quad (7)$$

where the first line is due to the convexity of an exponential function and Jensen's inequality:

$$\begin{aligned} & \exp \left(\mathbb{E}_{\mathbb{P}^\infty} \left[\max_{t \in [T]} (f_t(\mathbf{x}, \tilde{\mathbf{z}}) - f_t(\mathbf{x}, \boldsymbol{\lambda})) / \mu \right] \right) \\ &\leq \mathbb{E}_{\mathbb{P}^\infty} \left[\exp \left(\max_{t \in [T]} (f_t(\mathbf{x}, \tilde{\mathbf{z}}) - f_t(\mathbf{x}, \boldsymbol{\lambda})) / \mu \right) \right]; \end{aligned}$$

the fourth line results from substituting $f_t(\mathbf{x}, \tilde{\mathbf{z}}) = \sum_{v \in \mathcal{V}_t} x_{v,t} \tilde{z}_v$ into the right-hand side (RHS) of the third line; and the fifth and sixth lines follow from (4).

If we combine (6) and (7) as follows:

$$\begin{aligned} & \mathbb{E}_{\mathbb{P}^\infty} \left[\max_{t \in [T]} f_t(\mathbf{x}, \tilde{\mathbf{z}}) \right] \\ &\leq \inf_{\mu > 0} \left(\mu \ln \sum_{t \in [T]} \exp \left(\sum_{v \in \mathcal{V}_t} \lambda_v (e^{x_{v,t}/\mu} - 1 - x_{v,t}/\mu) \right) \right) \\ &\quad + \max_{t \in [T]} f_t(\mathbf{x}, \boldsymbol{\lambda}), \end{aligned}$$

then its RHS has the following epigraph form

$$\begin{aligned} & \inf_{\mu > 0, \kappa, \gamma} \kappa + \gamma \\ & \text{s.t. } \sum_{v \in \mathcal{V}_t} x_{v,t} \lambda_v \leq \gamma, \quad \forall t \in [T] \\ & \mu \ln \sum_{t \in [T]} \exp \left(\sum_{v \in \mathcal{V}_t} \lambda_v (e^{x_{v,t}/\mu} - 1 - x_{v,t}/\mu) \right) \leq \kappa. \end{aligned} \quad (8)$$

Constraint (8) is equivalent to $\sum_{t \in [T]} \mu \exp(\sum_{v \in \mathcal{V}_t} \lambda_v (e^{x_{v,t}/\mu} - 1 - x_{v,t}/\mu) - \kappa/\mu) \leq \mu$, where the expression within the first summation of the left-hand side, that is, $\mu \exp(\sum_{v \in \mathcal{V}_t} \lambda_v (e^{x_{v,t}/\mu} - 1 - x_{v,t}/\mu) - \kappa/\mu)$, has an epigraph form for all $t \in [T]$ as follows:

$$\begin{aligned} & \inf_{\zeta_t} \zeta_t \\ & \text{s.t. } \mu \exp \left(\left(-\kappa + \sum_{v \in \mathcal{V}_t} \lambda_v (\mu e^{x_{v,t}/\mu} - x_{v,t} - \mu) \right) / \mu \right) \leq \zeta_t. \end{aligned} \quad (9)$$

Furthermore, the term $\mu e^{x_{v,t}/\mu}$ inside Constraint (9) has an epigraph form $\inf_{\xi_{v,t}} \{\xi_{v,t} : \mu e^{x_{v,t}/\mu} \leq \xi_{v,t}\}$ for all $v \in \mathcal{V}_t$. Hence, the epigraph containing Constraint (8) can be equivalently written as

$$\begin{aligned} & \inf_{\mu > 0, \kappa, \gamma, \xi, \zeta} \kappa + \gamma \\ & \text{s.t. } (5a) - (5d). \end{aligned} \quad (10)$$

Because $f_s(\mathbf{x}, \tilde{\mathbf{z}})$ is linear in $\tilde{\mathbf{z}}$ for all $s \in [T]$ and $\tilde{\mathbf{z}} \sim \mathbb{P}^\infty$ is the vector of independent Poisson random variables, we have

$$\mathbb{E}_{\mathbb{P}^\infty} [f_s(\mathbf{x}, \tilde{\mathbf{z}})] = f_s(\mathbf{x}, \boldsymbol{\lambda}). \quad (11)$$

Therefore, by combining (6), (7), (10), and (11) and optimizing over x , we get the optimal value of ECP-U as an upper bound of π^* , as a solution optimal to ECP-U is also feasible to (3). $\square$

The model in Proposition 1 is labeled ECP-U (recall U represents “uncapacitated”) because all the constraints in this model involve either linear or exponential functions and thus can be expressed as exponential cone constraints. For example, Constraint (5b), $\mu \exp(x_{v,t}/\mu) \leq \xi_{v,t}$, can be written as $(\xi_{v,t}, \mu, x_{v,t}) \in \mathcal{K}_{\text{exp}}$. Hence, ECP-U is an ECP, and thus can be solved via a state-of-the-art conic programming solver efficiently (MOSEK ApS 2020). The number of decision variables and constraints of ECP-U are on the same order as those of our SP model (3).

Remark 1. To the best of our knowledge, the only other paper to approximate an SP with ECPs is See and Sim (2010), but their approach requires the random variables of the SP to be sub-Gaussian to derive meaningful bounds, which is not applicable to our SP. Moreover, our approximation method differs from theirs by directly bounding the expectation of piecewise linear convex functions while they first use a linear decision rule and then approximate the resulting expected surplus of affine functions. Finally, we directly solve ECPs, whereas all previous papers (e.g., See and Sim 2010) approximate ECPs using second-order cone programs.

Remark 2. Although developed for a specific charging problem, our ECP approach of bounding the expected recourse function could be potentially useful for solving more general two-stage SPs. Consider a two-stage linear SP with fixed recourse where the recourse function is defined generally as

$$Q(x, \tilde{z}) = \inf_y \{d'y : H(\tilde{z})x + Wy \geq h(\tilde{z})\},$$

where $h(\tilde{z})$ and $H(\tilde{z})$ are affine in $\tilde{z}$. If the dual feasible set $\{q \geq 0 : W'q = d\}$ is nonempty and compact with vertices v_i , $i \in [I]$, the expected recourse function becomes

$$\mathbb{E}_{\mathbb{P}}[Q(x, \tilde{z})] = \mathbb{E}_{\mathbb{P}} \left[\max_{i \in [I]} v_i' (h(\tilde{z}) - H(\tilde{z})x) \right].$$

If $\tilde{z}$ has independent components with exponential cone representable log-MGF, this expected recourse function can be upper bounded by an ECP using the same approach as in Proposition 1.

We denote an optimal solution to ECP-U by $\bar{x}$ and the value of the objective function in (3) evaluated at $\bar{x}$ by π^E , that is, $\pi^E = \mathbb{E}_{\mathbb{P}^\infty}[c(\bar{x}, \tilde{z})]$. We next show that compared with an optimal solution to (3), x^* , there is a theoretical guarantee on the performance of $\bar{x}$. Theorem 1 compares π^E and π^* ($\pi^* = \mathbb{E}_{\mathbb{P}^\infty}[c(x^*, \tilde{z})] = \mathbb{E}_{\mathbb{P}^C}[c(x^*, \tilde{z})]$, as $\mathbb{P}^C \equiv \mathbb{P}^\infty$ when C is infinite).

Theorem 1. Any optimal solution $\bar{x}$ to ECP-U has a performance guarantee as follows:

$$\begin{aligned} \pi^E - \pi^* &= \mathbb{E}_{\mathbb{P}^\infty}[c(\bar{x}, \tilde{z})] - \mathbb{E}_{\mathbb{P}^\infty}[c(x^*, \tilde{z})] \\ &\leq d \inf_{\mu > 0} \left(\mu \ln \sum_{t \in [T]} \exp \left(\sum_{v \in \mathcal{V}_t} \lambda_v \phi \left(\frac{\min\{u_v, K\}}{\mu \eta} \right) \right) \right), \end{aligned} \quad (12)$$

where $\phi(x) \triangleq e^x - 1 - x$.

Theorem 1 gives in (12) an upper bound on the difference between π^E and π^* . This bound depends on d , η , K , and T , as well as λ_v and u_v (for all $v \in [V]$), but it does not depend on e_s . Although this bound is not tight, it can be computed *before* solving ECP-U. *After* solving ECP-U, the bound on the difference between π^E and π^* on our numerical instances becomes tighter (see Section 7).

4.2. Capacitated Case

Unlike the uncapacitated case, $\tilde{z} \sim \mathbb{P}^C$ does not follow independent Poisson distributions when $C < \infty$, so it is difficult to obtain its analytical characterization analogous to that in (4). We then use the idea of DRO: It is often used to mitigate distributional ambiguity, but we use DRO to derive an ECP approximation as a tractable upper bound of our SP model. In our SP model, the arrival distribution is *known* (joint truncated Poisson distribution, $\tilde{z} \sim \mathbb{P}^C$), and thus there is no distributional ambiguity. We relax this known distribution to be contained in an ambiguity set in the DRO framework, which we denote by $\mathcal{F}$, that is, $\mathbb{P}^C \in \mathcal{F}$. Therefore, the solution from the DRO framework gives an upper bound of our SP model with the known distribution. As we show later, the DRO formulation is intractable, and thus we obtain an upper bound of this DRO formulation by developing an ECP approximation, which in turn gives an upper bound of our SP model.

We use the infinitely constrained “entropic dominance” ambiguity set, adapted from Chen et al. (2019):

$$\mathcal{F} \triangleq \left\{ \mathbb{P} \in \mathcal{P}_0(\mathbb{R}^V) \left| \begin{array}{l} \tilde{z} \sim \mathbb{P} \\ \ln \mathbb{E}_{\mathbb{P}}[\exp(\theta' \tilde{z})] \leq \sum_{v \in [V]} \lambda_v (e^{\theta_v} - 1), \\ \forall \theta \geq 0 \\ \mathbb{P}[\tilde{z} \in \mathcal{Z}] = 1 \end{array} \right. \right\},$$

where the third line is the support constraint with the support $\mathcal{Z}$ given in (1). The second line is entropic dominance constraints, which restrict *all* the moments of random variables $\tilde{z}_v$'s to be no greater than those of independent Poisson variables (i.e., in uncapacitated case): The left- and right-hand sides of this set of

constraints are the log-MGFs of $\tilde{z} \sim \mathbb{P}^C$ and $\tilde{z} \sim \mathbb{P}^\infty$ (see (4)), respectively. To further understand $\mathcal{F}$, we write it as the intersection of the following two sets, as shown in Lemma 1:

$$\begin{aligned}\mathcal{F}^1 &\triangleq \left\{ \mathbb{P} \in \mathcal{P}_0(\mathbb{R}^V) \mid \tilde{z} \sim \mathbb{P}; \ln \mathbb{E}_{\mathbb{P}}[\exp(\boldsymbol{\theta}' \tilde{z})] \right. \\ &\quad \left. \leq \sum_{v \in [V]} \lambda_v (e^{\theta_v} - 1), \forall \boldsymbol{\theta} \geq \mathbf{0} \right\} \\ \mathcal{F}^2 &\triangleq \{ \mathbb{P} \in \mathcal{P}_0(\mathbb{R}^V) \mid \tilde{z} \sim \mathbb{P}; \mathbb{E}_{\mathbb{P}}[\tilde{z}] \leq \boldsymbol{\lambda}; \mathbb{P}[\tilde{z} \in \mathcal{Z}] = 1 \}.\end{aligned}$$

Lemma 1. The ambiguity set $\mathcal{F} = \mathcal{F}^1 \cap \mathcal{F}^2$.

Lemma 1 shows the following. (1) Both the entropic dominance constraints (in $\mathcal{F}^1$) and the support constraint (in $\mathcal{F}^2$) are crucial in defining the ambiguity set. For instance, we show numerically that a model without $\mathcal{F}^1$ performs much worse in Section 7.4.4. (2) As $\mathcal{F}^2$ is an example of the moment-based ambiguity sets commonly used in DRO that involves only a constraint on the first moment (i.e., $\mathbb{E}_{\mathbb{P}}[\tilde{z}] \leq \boldsymbol{\lambda}$), Lemma 1 implies that $\mathcal{F}$, as a subset of $\mathcal{F}^2$, incorporates more moment information than $\mathcal{F}^2$ does. Therefore, $\mathcal{F}$ is smaller and less conservative than $\mathcal{F}^2$. Next, we show that the ambiguity set $\mathcal{F}$ contains the given distribution $\mathbb{P}^C$.

Proposition 2. The probability distribution $\mathbb{P}^C \in \mathcal{F}$.

Therefore, we can obtain an upper bound of the optimal value of (3) by considering the worst-case expected total cost over the ambiguity set $\mathcal{F}$:

$$\min_{x \in \mathcal{X}} \sup_{\mathbb{P} \in \mathcal{F}} \mathbb{E}_{\mathbb{P}}[c(x, \tilde{z})], \quad (\text{DRO-Ent})$$

where “Ent” represents “entropic.” However, because of the infinite number of constraints in $\mathcal{F}$ (as the entropic dominance constraints have to hold for all $\boldsymbol{\theta} \geq \mathbf{0}$), the DRO-Ent model is not tractable (Chen et al. 2019). Therefore, we approximate DRO-Ent using an ECP formulation similar to that in Proposition 1, which can also be solved efficiently. We use Lemma 1 to write DRO-Ent as the optimization over ambiguity set $\mathcal{F}^1$ and $\mathcal{F}^2$ separately. For the former case, we can use a technique similar to that of the proof of Proposition 1; for the latter, as $\mathcal{F}^2$ is an example of moment-based ambiguity sets used in DRO, we use techniques standard in DRO to obtain bounds. We combine these two to obtain a tractable upper bound as follows, which we label as ECP-C (where C represents “capacitated” as mentioned previously).

Proposition 3. When $C < \infty$, the following ECP-C gives an upper bound of π^* :

$$\inf_{\substack{x \in \mathcal{X}, a, b \geq 0, v \geq 0, y \geq 0, \mu > 0, \\ \kappa, \gamma, \alpha, \beta \geq 0, \xi, \zeta, \rho \geq 0}} d(\kappa + \gamma + \alpha + \beta' \boldsymbol{\lambda}) + a + b' \boldsymbol{\lambda}$$

$$\text{s.t. } \sum_{v \in \mathcal{V}_t} y_{v,t} \lambda_v \leq \gamma \quad \forall t \in [T], \quad (13a)$$

$$(\xi_{v,t}, \mu, y_{v,t}) \in \mathcal{K}_{\text{exp}} \quad \forall t \in [T], v \in \mathcal{V}_t, \quad (13b)$$

$$\left(\zeta_t, \mu, -\kappa + \sum_{v \in \mathcal{V}_t} \lambda_v (\xi_{v,t} - y_{v,t} - \mu) \right) \in \mathcal{K}_{\text{exp}} \quad \forall t \in [T], \quad (13c)$$

$$\sum_{t \in [T]} \zeta_t \leq \mu, \quad (13d)$$

$$C \sum_{k \in [T]} \rho_t^k \leq \alpha \quad \forall t \in [T], \quad (13e)$$

$$x_{v,t} - y_{v,t} - \beta_v \leq \sum_{k \in \mathcal{T}_v} \rho_t^k \quad \forall t \in [T], v \in \mathcal{V}_t, \quad (13f)$$

$$C \sum_{t \in [T]} v_t \leq a, \quad (13g)$$

$$\sum_{s \in \mathcal{T}_v} x_{v,s} e_s - b_v \leq \sum_{t \in \mathcal{T}_v} v_t \quad \forall v \in [V]. \quad (13h)$$

Similar to ECP-U, ECP-C can be solved efficiently. (The number of decision variables and constraints in ECP-C is on the same order as those in our SP model (3).)

Theorem 2. The optimal value of ECP-C increases as C increases and converges to a value that is upper bounded by the optimal value of ECP-U as $C \rightarrow \infty$.

Theorem 2 shows that the optimal value of ECP-C increases as charging station capacity increases, intuitively because more vehicles will be admitted to the station to be charged. It also shows that this optimal value of ECP-C converges and yields an asymptotic lower bound of the optimal value of ECP-U. However, unlike ECP-U, we are unable to obtain a theoretical performance guarantee of an optimal action to ECP-C and thus use a numerical study to examine its performance together with those of two benchmark approaches.

5. Benchmark Approximations

5.1. SAA

The most common approach to solve a two-stage SP is SAA based on Monte Carlo sampling: Generate from the given distribution many independent samples and replace this distribution in the SP model with the empirical distribution of these samples. We denote these samples from $\mathbb{P}^C$ by $z^i \equiv (z_v^i)_{v \in [V]} \forall i \in [N]$, where N is the sample size. The SAA model can be obtained after reformulating (3) as a two-stage SP using auxiliary

variables:

$$\begin{aligned} \inf_{x \in \mathcal{X}, \gamma^i \forall i \in [N]} & \frac{1}{N} \sum_{i \in [N]} \left(d\gamma^i + \sum_{s \in [T]} e_s \left(\sum_{v \in \mathcal{V}_s} x_{v,s} z_v^i \right) \right) \\ \text{s.t.} & \sum_{v \in \mathcal{V}_i} x_{v,t} z_v^i \leq \gamma^i \quad \forall t \in [T], i \in [N]. \end{aligned} \quad (14)$$

Let us denote by x^S an optimal solution to (14) and by π^S the value of the objective function in (3) evaluated at x^S , that is, $\pi^S = \mathbb{E}_{\mathbb{P}^C}[c(x^S, \tilde{z})]$. Because x^S and π^S depend on the samples generated and are thus random, to obtain a solution near optimal to (3) with a high probability requires the sample size N to be sufficiently large, especially given the large-scale nature of (3). In particular, as mentioned in Shapiro and Nemirovski (2005), the sample complexity of obtaining an ϵ -optimal solution is $\tilde{O}\left(\frac{D^2 L^2}{\epsilon^2}\right)$, where L is the Lipschitz constant of the objective function $f(\cdot, z)$ for all $z \in \mathcal{Z}$ and D is the diameter of $\mathcal{X}$. Nonetheless, fewer samples may be needed in practice. For instance, in our numerical study, we need about 38,000 samples to obtain a near-optimal solution (see Figure 3).

5.2. DRO

Another benchmark is a DRO approach with the mean-variance ambiguity set:

$$\min_{x \in \mathcal{X}} \sup_{\mathbb{P} \in \mathcal{F}_\Sigma} \mathbb{E}_{\mathbb{P}}[c(x, \tilde{z})], \quad (\text{DRO-C})$$

where

$$\mathcal{F}_\Sigma = \left\{ \mathbb{P} \in \mathcal{P}_0(\mathbb{R}^V) \left| \begin{array}{l} \tilde{z} \sim \mathbb{P} \\ \mathbb{E}_{\mathbb{P}}[\tilde{z}] = \mu \\ \mathbb{E}_{\mathbb{P}}[\tilde{z}\tilde{z}'] \preceq \Sigma + \mu\mu' \\ \mathbb{P}[\tilde{z} \in \mathcal{Z}] = 1 \end{array} \right. \right\}. \quad (15)$$

Using standard techniques in DRO (e.g., Wiesemann et al. 2014), the model DRO-C can be reformulated as the following SDP (see Online Appendix A for its derivation):

$$\begin{aligned} \min_{\substack{x \in \mathcal{X}, \alpha, \beta, \Gamma, \\ \xi^t \geq 0, \zeta^t \geq 0, y^t, \forall t \in [T]}} & d(\alpha + \beta' \mu + \Gamma \bullet (\mu\mu' + \Sigma)) \\ & + \sum_{s \in [T]} e_s f_s(x, \mu) \\ \text{s.t.} & \begin{pmatrix} \alpha - C \sum_{s \in [T]} \zeta_s^t & y^t \\ y^t & \Gamma \end{pmatrix} \succeq 0 \quad \forall t \in [T] \\ & 2y_v^t = \beta_v - x_{v,t} - \xi_v^t + \sum_{s \in \mathcal{I}_v} \zeta_s^t \\ & \quad \forall t \in [T], \forall v \in \mathcal{V}_t \\ & 2y_v^t = \beta_v - \xi_v^t + \sum_{s \in \mathcal{I}_v} \zeta_s^t \\ & \quad \forall t \in [T], \forall v \in [V] \setminus \mathcal{V}_t. \end{aligned}$$

This SDP formulation has T semidefinite constraints of dimension $V+1$ and thus may run into computation time and memory issues when V and T are large. In general, the complexity of SDP is higher than ECP (which is the same as SOCP): Obtaining ϵ -optimal solutions to SDP can be found in $O(n^2 k^{2.5} \log(1/\epsilon))$ arithmetic operations via interior point methods, where n is the number of decision variables and k is the dimension of the symmetric matrices in the SDP constraints; obtaining ϵ -optimal solutions to ECPs requires $O(n^2 k^{1.5} \log(1/\epsilon))$ arithmetic operations, where n is the number of decision variables and k is the number of exponential cone constraints (Nesterov and Nemirovskii 1994).

6. Joint Pricing and Scheduling

We consider an extension of the model in Section 3 to maximize the expected profit by jointly optimizing the pricing of charging service fees and scheduling of EVs. We specify the formulation of this extended model in Section 6.1. We solve this model by leveraging the ECP approximations from Section 4 and consider the uncapped and capacitated cases in Section 6.2 and Section 6.3, respectively.

6.1. Problem Formulation

Service Fee Structure. We use one of the most common pricing schemes in practice, where the service fee is proportional to the quantity of electricity charged (also known as kWh-based) (Blink 2020, Tesla 2020). This means that in our model customers of type v pay pu_v for the EV charging service, where $p \in [\underline{p}, \bar{p}]$ is the decision of unit service fee.

Price-Dependent Arrival Rate. For ease of exposition and computational tractability, we assume that the customer arrival rate is linearly decreasing in the price:

$$\lambda_v = \bar{\lambda}_v(1 - r_v p), \quad \forall v \in [V], \quad (16)$$

where $\bar{\lambda}_v$ is the maximum possible arrival rate of type v customers (e.g., when the EV charging service is free), and $r_v > 0$ is the price sensitivity of customer type v . Consistent with the literature in the pricing of service operations (e.g., Mendelson and Whang 1990), this price-dependent arrival rate represents aggregating the individual customer's decision of whether to use the charging service according to their service valuations, which follow a uniform distribution $F(\cdot)$ between $[0, 1/r_v]$: Because only customers with valuation higher than p join the service at the charging station, the effective arrival rate of type v customers is given by $\lambda_v = \bar{\lambda}_v(1 - F(p)) = \bar{\lambda}_v(1 - r_v p)$.

Objective Function. The objective is to jointly optimize the pricing decision p and the scheduling decision x to maximize the expected total profit, that is, the total

revenue minus total cost, as follows:

$$\max_{p \in [\underline{p}, \bar{p}], x \in \mathcal{X}} \mathbb{E}_{\mathbb{P}^C} \left[\sum_{v \in [V]} pu_v \tilde{z}_v - c(x, \tilde{z}) \right]. \quad (17)$$

Problem (17) is much more challenging than the scheduling problem (3) because the underlying distribution $\mathbb{P}^C$ in (17) depends on the pricing decision p . To solve (17) efficiently, we leverage our ECP approximations in Section 4 for both the uncapacitated and capacitated cases.

6.2. Uncapacitated Case

Proposition 4. When $C = \infty$, the optimal value of JPS-U is a lower bound of (17):

$$\begin{aligned} & \sup_{x \in \mathcal{X}, p \in [\underline{p}, \bar{p}], \mu > 0, \lambda, \kappa, \gamma} \sum_{v \in [V]} pu_v \lambda_v - d(\kappa + \gamma) \\ & \quad - \sum_{s \in [T]} e_s f_s(x, \lambda) \quad (\text{JPS-U}) \\ & \text{s.t.} \quad (5a), (8), (16). \end{aligned}$$

The ECP formulation in Proposition 4 is labeled JPS-U, representing “joint pricing and scheduling” for the “uncapacitated” case. Unlike ECP-U, which is convex, JPS-U is nonconvex as (8) defining JPS-U is nonconvex because p (and equivalently λ_v) is a decision variable. Therefore, we propose an optimization procedure to solve JPS-U efficiently by alternating between fixing p and fixing x and μ : Both of the resulting formulations are convex and can be easily solved using off-the-shelf solvers. Specifically, first, given p , the arrival rates λ_v are given in (16) and thus fixed, so we can follow Proposition 1 to reformulate JPS-U as the following convex optimization problem with only scheduling decisions:

$$\begin{aligned} & \sup_{x \in \mathcal{X}, \mu > 0, \kappa, \gamma, \xi, \zeta} \sum_{v \in [V]} pu_v \lambda_v - d(\kappa + \gamma) - \sum_{s \in [T]} e_s f_s(x, \lambda) \\ & \text{s.t.} \quad (5a) - (5d). \end{aligned} \quad (18)$$

Second, given x and μ , JPS-U can also be reformulated as the following convex optimization problem as in Proposition 1:

$$\begin{aligned} & \sup_{p \in [\underline{p}, \bar{p}], \lambda, \kappa, \gamma, \xi, \zeta} \sum_{v \in [V]} u_v \bar{\lambda}_v (p - r_v p^2) - d(\kappa + \gamma) \\ & \quad - \sum_{s \in [T]} e_s f_s(x, \lambda) \quad (19) \\ & \text{s.t.} \quad (5a), (5d), (9), (16). \end{aligned}$$

Therefore, we can solve JPS-U by alternatingly solving Problems (18) and (19) with the following algorithm, the convergence of which is shown in Proposition 5.

Algorithm 1 (Alternating Optimization for JPS-U)

1. **Initialization** Set initial price $p^{(0)} \in [\underline{p}, \bar{p}]$, iteration counter $i \leftarrow 1$
2. **Scheduling optimization** Solve Model (18) with input $p^{(i-1)}$ and let x^* and μ^* be the optimal solution; set $x^{(i-1)} \leftarrow x^*$ and $\mu^{(i-1)} \leftarrow \mu^*$. Store optimal value $val_1^{(i)}$;
3. **Pricing optimization** Solve Model (19) with inputs $x^{(i-1)}$ and $\mu^{(i-1)}$, and let p^* be the optimal solution; set $p^{(i)} \leftarrow p^*$. Store optimal value $val_2^{(i)}$;
4. **Termination** If $|val_2^{(i)} - val_1^{(i)}| < \delta$ (where $\delta > 0$ is a given small tolerance), set $p^* \leftarrow p^{(i)}$, solve Model (18) with input $p^{(i)}$ to obtain x^* , output p^* and x^* , and then stop. Otherwise, set $i \leftarrow i + 1$ and go back to Step 2.

Output: Pricing decision p^* and scheduling decision x^*

Proposition 5. The sequence of optimal values $\{val_1^{(i)}, val_2^{(i)}\}$ in Algorithm 1 is nondecreasing and converges to a finite value.

6.3. Capacitated Case

Next, we develop an ECP approximation of Model (17) for the capacitated case. We first derive a convex upper bound for the expected revenue in Proposition 6. Then this upper bound on the expected revenue minus the upper bound on the expected cost in Proposition 3 gives an approximation of the expected profit in (17). In the capacitated case, there is only an approximation, different from the uncapacitated case where we obtain a lower bound.

Proposition 6. The optimal value of the following convex optimization problem is an upper bound of the expected revenue $\mathbb{E}_{\mathbb{P}^C}[\sum_{v \in [V]} pu_v \tilde{z}_v]$:

$$\begin{aligned} & \sup_{g \geq 0} \sum_{v \in [V]} g_v u_v \\ & \text{s.t.} \quad \bar{\lambda}_v (p - r_v p^2) \geq g_v \quad \forall v \in [V], \end{aligned} \quad (20a)$$

$$Cp \geq \sum_{v \in \mathcal{V}_t} g_v \quad \forall t \in [T]. \quad (20b)$$

Proof. Suppose $\mathbb{E}_{\mathbb{P}^C}[\tilde{z}] = \underline{\lambda}$, then $\mathbb{E}_{\mathbb{P}^C}[\sum_{v \in [V]} pu_v \tilde{z}_v] = \sum_{v \in [V]} pu_v \underline{\lambda}_v$. We know $\underline{\lambda} \leq \lambda$ because of arrival truncation in capacitated case. The support constraint $\mathbb{P}^C[\tilde{z} \in \mathcal{Z}] = 1$ implies $\sum_{v \in \mathcal{V}_t} \underline{\lambda}_v \leq C$, $\forall t \in [T]$, from (1). Hence,

$$\sup_{b \geq 0} \left\{ \sum_{v \in [V]} pu_v b_v : b \leq \lambda, \sum_{v \in \mathcal{V}_t} b_v \leq C, \forall t \in [T] \right\} \quad (21)$$

is an upper bound of $\sum_{v \in [V]} pu_v \underline{\lambda}_v$ as $\underline{\lambda}$ is feasible in Problem (21). Next, we perform a change of variable by

$g_v = pb_v$. As $p > 0$, Problem (21) is equivalent to

$$\sup_{g \geq 0} \left\{ \sum_{v \in [V]} u_v g_v : g \leq \lambda p, \sum_{v \in V_t} g_v \leq Cp, \forall t \in [T] \right\}. \quad (22)$$

Substituting (16) into Model (22) concludes the proof. $\square$

Combining Propositions 6 and 3, which give tractable approximations of the expected revenue and cost, respectively, we obtain the following approximation of the expected profit:

$$\begin{aligned} \sup_{\substack{x \in \mathcal{X}, p \in [p, \bar{p}], \lambda, a, b \geq 0, \nu \geq 0, y \geq 0, \\ \mu > 0, \kappa, \gamma, \alpha, \beta \geq 0, \xi, \zeta, \rho \geq 0, g \geq 0}} \sum_{v \in [V]} g_v u_v - d(\kappa + \gamma + \alpha + \beta' \lambda) \\ - a - b' \lambda \end{aligned} \quad (\text{JPS-C})$$

s.t. (13a) – (13h), (16), (20a) – (20b).

As in the uncapacitated case, we can perform alternating optimization to efficiently solve problem JPS-C using an algorithm analogous to Algorithm 1: Fixing p (or equivalently λ) in JPS-C, we obtain the scheduling optimization, which is an ECP; fixing all variables except $p, g, \lambda, \gamma, \kappa$ in JPS-C, we obtain the pricing optimization, which is a conic program involving exponential cones and second-order cones. This algorithm has a similar convergence property as Algorithm 1 and is omitted for brevity.

7. Numerical Study

In this section, we examine the performance of our ECP approximation in benchmark with the SAA and DRO approaches using numerical instances calibrated to real data for both the main model and the extended model with the pricing decision. We also discuss how they perform considering several practical implementation issues. Last, we use ECP to generate managerial insights.

7.1. Experiment Settings

We consider the planning horizon of a service provider to be one day as there are very few EVs charging across two consecutive days in public charging stations, as shown in our data. Each period represents 15 minutes, which is used in practice to compute demand charges for most utility companies, such as Southern California Edison. Therefore, the number of periods T equals 96.

Electricity Tariff Calibration. We use Southern California Edison's electricity tariff *Schedule GS-2 General Service* (Neubauer and Simpson 2015), which is the tariff an electricity user such as an EV charging service provider is subject to. Under this tariff, we obtain the unit energy charge and demand charge depending on the hour of


the day as follows:

$$\hat{e}_t = \begin{cases} \$0.1466/\text{kWh} & \text{if } 13 \leq [t/4] \leq 18 \text{ (on-peak hours)} \\ \$0.0895/\text{kWh} & \text{if } 9 \leq [t/4] \leq 12 \text{ or } 19 \leq [t/4] \leq 23 \\ & \text{(mid-peak hours)} \\ \$0.0582/\text{kWh} & \text{otherwise (off-peak hours),} \end{cases}$$

$$\hat{d}_t = \begin{cases} \$0.465/\text{kW} & \forall t \in [T] \doteq [96] \text{ (all-period)} \\ \$0.540/\text{kW} & \text{if } 13 \leq [t/4] \leq 18 \text{ (on-peak hours)} \\ \$0.165/\text{kW} & \text{if } 9 \leq [t/4] \leq 12 \text{ or } 19 \leq [t/4] \leq 23 \\ & \text{(mid-peak hours),} \end{cases}$$

where $[t/4]$ represents the hour index of period t (as one hour has four periods). These demand charge parameters for the horizon of a day are obtained by dividing those for a month by 30 to maintain the same proportion of the demand charge relative to the total cost. More details on the calibration can be found in Online Appendix B. The demand charge is the sum of all-period demand charge (applicable to all periods), on-peak demand charge, and mid-peak demand charge: Although the models in Sections 4 and 6 have only all-period demand charge for the expositional simplicity, the ECP models for the case with multiple demand charges are constructed analogously to those in Sections 4 and 6.

EV Arrival Rate Calibration. In this study, we use the public EV charging data from the UK Department for Transportation. Although we acknowledge the discrepancies in EV charging behaviors across different markets, we use the UK charging events to approximate general EV arrival patterns in the computational experiments. The dataset in the United Kingdom contains all of the charging events from 27 local authorities in England from January 1 to December 31, 2017 (UK Department for Transport 2018). For each charging event, it provides the start time, end time, and electricity supplied. In particular, we consider the data of both rapid and fast chargers where the charging speed are above 22 kW and below 22 kW, respectively. The average charging duration for rapid chargers is short, about 35 minutes, and for fast chargers is long, about 4.7 hours. Although the value of using our model is higher for long-duration charging than for short-duration charging as the former gives more room for optimization, we focus on the latter instead due to the following reason: Although ECP scales well and can be solved for both the short- and long-duration charging, SAA does not scale well and cannot be solved to a satisfactory level for long-duration cases. Hence, unless specified explicitly otherwise, all the results mentioned later pertain to the short-duration charging case. We demonstrate that the value of our model remains high even for short-duration charging.

Figure 1. (Color online) Normalized Histogram of Customer Stay Duration and Electricity Charged at Rapid Chargers in 2017 (UK Department for Transportation 2018)

Limited experiments show that this value is much higher in the case of long-duration charging. We observe that the average charging speeds (electricity supplied divided by the stay duration, the time between arrival and departure) for most charging events are *not* at the maximum. This observation holds at the aggregate level across all locations, at the individual charging location (local authority) level, and at the individual charger level. This observation corroborates the need to optimize the scheduling of EV charging with various customer types, as studied in this paper. In the rest of the paper, we use the arrival pattern aggregated across all locations unless specified otherwise.


After preprocessing the dataset (see details in Online Appendix B), we obtain an average of 278.65 daily charging events. We then discretize time into 15-minute intervals. In particular, we count any partial arrival/departure period as a full period: For example, a car that arrives at 12:56 p.m. and leaves at 1:22 p.m. is treated as staying between 12:45 p.m. to 1:30 p.m.; that is, it stays for three 15-minute periods.⁵ Then we compute all policies based on this 15-minute discretization. We can easily convert this 15-minute policy to a policy that is implementable on a 1-minute discretization using a simple translation algorithm (see more details in Section 7.4.1). We show that the performance of this converted policy is close to that of the policy computed on the 1-minute discretization, and the performance comparison between the converted ECP and converted SAA is the same as that on 15-minute discretization. In addition, there are two more reasons to use 15-minute discretization: First, it is consistent with the measurement for the demand charge used by utility firms in the United States and thus often used in the literature (Glassmire et al. 2012, Zhang et al. 2017, Zhang and Augenbroe 2018). Second, finer time discretization levels are computationally more challenging due to memory and CPU time limitations; for instance, the number of random variables and

decision variables under 1-minute discretization is 15^2 and 15^3 , respectively, times of those under 15-minute discretization.

We plot in Figure 1 the resulting density (or normalized frequency) of EV stay duration (time between arrival and departure) and electricity charged. The mean (standard deviation) of stay duration is 35.83 (27.92) minutes; the mean (standard deviation) of electricity charged is 10.69 (7.64) kWh. We then estimate the arrival rate of each type as the average number of charging events of *each type* per day within the entire horizon. As our data span a year, *each type* has 365 data points for estimating its arrival rate: For each day, each type has *one realization* of its arrivals, even if the number of arrivals on that day is zero. Subsequently, among about 80,000 customer types with distinct start times, departure times, and total electricity charged, we identify 14,284 customer types with positive arrival rates. We use these arrival rates for the *given* truncated Poisson distributions in our SP model to compare the performances of all policies. Although we assume Poisson arrivals with known rates, there may be estimation errors in arrival rates due to insufficient data and model mis-specification of the arrival distribution in practice, and we therefore consider these two cases in Section 7.4.2.

Next, we consider managing a hypothetical charging station to charge EVs. We plot in Figure 2 the expected number of EVs at the charging station in a one-day horizon under various values of the number of chargers C to achieve different service levels (the probability that a customer finds a charger upon arrival). As seen, the expected number of customers at the station follows a daily pattern: It starts low from midnight to 6 a.m., increases significantly to the peak at around 1 p.m., and finally decreases significantly until midnight. When $C = 30$, there is almost no lost customer (with service level 99.86%): The corresponding line in Figure 2

Figure 2. (Color Online) Average Number of EVs at the Charging Station



almost coincides with that for the uncapacitated case. However, with $C = 25, 20$, and 15 , the service level drops quickly to 98.83%, 94.07%, and 82.54%, respectively.

Other Parameters. We use the specifications of Nissan Leaf, a popular EV type. Specifically, we set $\hat{U} = 62$ kWh using the energy capacity of Nissan Leaf (Nissan 2019); $\hat{K} = 10.75$ kWh/period based on the rapid AC chargers with power at 43 kW (Zap-Map 2019); and $\hat{\eta} = 0.9$ using the efficiency of lithium-ion battery in Nissan Leaf (Karlsson and Kushnir 2013).

Our 96-period instance is large-scale: There are in total 80,736 random variables and 692,702 decision variables. In this case, the benchmark DRO-C cannot be solved, and thus we focus on comparing ECP with SAA in the rest of this section. We also consider a numerical setup with small problem instances where there are only six periods and DRO-C can be solved. See Online Appendix C for more details of the comparisons among all three approaches.

7.2. Computation

We implement SAA using the following two methods: Solving the extensive form by the interior point (also known as barrier) method and the L-shaped method (also known as Benders decomposition) with either single cut or multi-cut (Birge and Louveaux 2011). In the L-shaped method, we first generate optimality cuts (as our problem has complete recourse) by solving N sub-problems over a simplex and then add these cuts to the master problem. In our multi-cut implementation, we do not add generated cuts for all the scenarios in each iteration, as the master problem (where the computational burden lies) is so large that we can only perform four iterations within 12 hours (our specified maximum computation time). Instead, we add into the master problem in each iteration the 100 cuts that have the largest gaps between the recourse function and its piecewise

affine lower bound at the current solution. We use dual simplex without presolve to solve all the linear programming problems in the L-shaped method.

As in our SP model the arrival distribution is *given*, we compute and evaluate ECP and SAA using the same calibrated arrival distribution. In particular, we generate 8,000 sample paths to compute the actions x^S under SAA. (Computing the action x^E under ECP does not require samples; only *evaluating* the action under ECP does.) After solving for x^E and x^S using the calibrated arrival distributions and other parameters, we generate from the same distribution via Monte Carlo simulation 10,000 samples to evaluate both, sufficient to show statistically significant performance comparison between both approaches. (The results do not change when using more samples in the evaluation step.) We evaluate x^E and x^S on these samples by computing the mean and standard deviation of $c(x^E, \tilde{z})$ and $c(x^S, \tilde{z})$, respectively. We implement SAA five times as the solutions obtained from SAA are sample dependent and report the final performance as the average of these five instances.


All of our numerical experiments are conducted in Python while using the MOSEK 9.2 solver on a Windows-OS 64-bit with 24 GB of RAM and an Intel i7-8750H CPU@2.20 GHz processor. All the problem instances except the L-shaped method are solved via the interior point method. For all methods, we set both primal and dual feasibility tolerance 10^{-8} and relative duality gap 10^{-8} . When the solver cannot find solutions to this accuracy, we accept the solution if it is feasible but not necessarily optimal.

7.3. Comparing ECP with SAA

Table 1 compares under different values of C the performances of ECP and SAA. In this table, SAA uses the interior point method, as it is faster than the L-shape method with either single cut or multi-cut (see such experiments on small six-period instances in Online Appendix C). As

Table 1. Performance Comparison Between ECP-C (or ECP-U) and SAA

C	Method	Total cost		CPU time (s)
		Mean	Standard deviation	
15	ECP-C	600.40	35.20	312.33
	SAA	608.91	36.76	13,015.66
20	ECP-C	717.73	44.42	354.66
	SAA	727.38	46.22	15,806.33
25	ECP-C	783.66	57.13	445.66
	SAA	794.04	59.44	16,207.17
30	ECP-C	805.00	67.55	320.56
	SAA	816.39	69.94	16,005.50
∞	ECP-U	808.90	71.30	194.22
	SAA	819.80	73.64	17,673.80

Figure 3. (Color Online) Expected Cost Under SAA at Different Sample Sizes Given $C = 30$ 

mentioned in Section 7.2, SAA uses 8,000 samples in computation, as SAA with a larger sample size cannot be solved within 12 hours (see more details on SAA with different sample sizes in Figure 11 in Online Appendix D). This table shows that compared with SAA, ECP results in not only a lower standard deviation of the mean cost but also a lower mean cost itself. (This difference in mean cost is significant compared with the standard error, which is the standard deviation of the total cost divided by the square root of the evaluation sample size of 10,000.) For instance, when C equals 30, ECP-C has a standard deviation and mean cost of 67.55 and 805, respectively, lower than those of SAA, which are 69.94 and 816.39, respectively. That is, the robustness of ECP-C (or lower standard deviation) does not come at the cost of a higher mean cost, contrary to known tradeoff (Choi and Ruszczyński 2008, Hao et al. 2020). Similar results are also observed when comparing ECP with

SAA using each arrival pattern from the three charging locations (local authorities) that have the largest number of charging events in the UK data.⁶ ECP-C beats SAA on robustness as expected because ECP-C leverages the ambiguity set inspired by DRO approaches (see Section 4.2). Compared with ECP models, SAA results in a high mean cost because it does not scale well and is thus hard to solve to optimality on large-scale instances.

To illustrate the reason why SAA has a higher mean cost than ECP, Figure 3 plots the extrapolation on the objective value of SAA (i.e., that of (14)) and the evaluation of SAA (i.e., π^S) at different sample sizes given $C = 30$: The number of samples required for π^S to be lower than π^E at $C = 30$ (which is 805, see Table 1) is large (about 23,000); that required for solving SAA to optimality is even larger (about 38,000), resulting in an estimated optimal mean cost of 799.30. However, this mean cost is not attainable because solving SAA to


Figure 4. (Color Online) Objective Value and Total CPU Time of Solving SAA with 38,000 Samples Using the L-Shaped Method with Multi-cut Given $C = 30$ 

Table 2. Performances of ECP-C and SAA on Long-Duration Instances

C	Method	Total cost		$\mathbb{E}\{\text{Demand charge}/\text{total cost}\}$
		Mean	Standard deviation	
60	ECP-C	404.33	31.52	0.58
	SAA(1,000)	454.71	38.31	0.64
50	ECP-C	372.67	27.93	0.58
	SAA(1,000)	417.55	34.53	0.64
40	ECP-C	327.68	25.78	0.58
	SAA(1,000)	350.92	29.93	0.62

Note. The number in the parentheses after SAA indicates the number of sample paths used to compute SAA.

near-optimality on our instances with 38,000 samples is computation-prohibitive, where we use the L-shaped method with multi-cut (as the interior method cannot be used due to memory issues). To demonstrate that, Figure 4 shows that SAA cannot converge even within 36 hours: Figure 4(a) plots both the lower bound of the objective value (obtained as the optimal value of the master problem) and the upper bound (obtained as the objective value of SAA at current solution), and Figure 4(b) shows that the CPU time per iteration increases exponentially when the number of iteration is large. The SAA solution using 38,000 samples at 36 hours leads to a mean cost of 866.06 and a standard deviation of 77.21, larger than those of SAA in Table 1.

Table 1 also lists the CPU computation time (excluding the evaluation time) of ECP and SAA in the last column. As seen, ECP computes about 50 times faster than SAA, suggesting that ECP is a more efficient computational approach to solve large-scale instances.

The superior performance of ECP over SAA is even more pronounced on numerical instances of long-duration charging, about 4.7 hours with fast chargers (see Section 7.1). Table 2 displays the mean and standard deviation of the total cost under both approaches, which shows that both measures under ECP-C are about 10% lower than those of SAA. This is mainly because SAA does not scale well, so we can compute SAA with only 1,000 sample paths due to memory limit: The number of customer types is much larger in the long-duration charging case, so we can no longer use the same number of sample paths (i.e., 8,000) for computing x^S from SAA as in the short-duration charging case. (Note that when evaluating both ECP-C and SAA, we continue using 10,000 sample paths from the arrival distributions calibrated for long-duration instances.) As such, we use short-duration charging cases in the rest of the paper.

We next show that ECP is near optimal as its optimality gap is at most around 4% in all instances. We estimate the optimality gap of ECP (i.e., $(\pi^E - \pi^*)/\pi^* \times 100\%$) using the optimal value of SAA (not listed in Table 1) as a lower bound of π^* (Shapiro et al. 2009). For instance, when $C = 30$, this optimality gap of ECP-C is at most $(805.00 -$

$773.67)/773.67 \times 100\% \approx 4.05\%$, where π^E is 805.00 (Table 1) and the optimal value of SAA given sample size 8,000 is 773.67. If we use the estimated optimal value 799.30 in Figure 3, the gap decreases to only 0.71%. We can also similarly estimate the optimality gap of ECP-U to be at most 4.24%. This gap is much smaller than that computed using the theoretical bound in (12), 214.67, divided by $\pi^E = 808.90$ (Table 1), which gives the optimality gap of 26.5%, as the former optimality gap is computed *after* solving ECP-U, whereas the latter is computed *before* solving it.

In summary, ECP performs better than other common approaches and achieves a near-optimal total cost. In Section 7.4, we consider different practical implementation issues in the numerical study. In Section 7.5, we also use ECP to generate managerial insights to demonstrate its practical relevance further.

7.4. Practical Considerations in Implementation

7.4.1. Time Discretization. As mentioned in Section 7.1, we discretize time into increments of 15 minutes and compute all policies on this discretization, and cannot discretize time into finer discretization, such as 1 minute, due to multiple reasons including CPU and memory limitations. We next examine the effect of our discretization by comparing the performance of our 15-minute policies with those on 1-minute discretization. To this end, we have to first convert our 15-minute policies to be implementable (i.e., feasible) on 1-minute discretization. If we simply divide the charging quantity for each type in each 15-minute period by the number of 1 minutes within this period, the resulting 1-minute level charging quantities in the arrival and departure periods are not necessarily feasible as they may violate the corresponding charging speed limit constraint in $\mathcal{X}$. For example, if the 15-minute policy in the arrival period is to charge at the maximum speed (i.e., the charging quantity is K/η), the 1-minute policy is to charge $K/(S\eta)$, where S is the number of 1-minute intervals in the arrival period: If $S < 15$, this exceeds $K/(15\eta)$, the charging speed limit under 1-minute discretization. Thus, before the division step, we have to shift the charging quantity on both the arrival and departure periods to the middle periods of the charging time window.

Algorithm 2 (Transform a 15-Minute Charging Policy $\mathbf{x}$ into a 1-Minute Implementable Policy $\mathbf{y}$)**Input:** 15-minute charging policy $x_{v,t}$

```

1 for 1-minute customer type  $v' \in [V']$  do
2   set  $s_v \leftarrow \lceil s_{v'}/15 \rceil$ ,  $\tau_v \leftarrow \lceil \tau_{v'}/15 \rceil$ ,  $v \leftarrow (s_v, \tau_v, u_{v'})$ ;
3   for 15-minute period  $t \in \{s_v, \dots, \tau_v\}$  do
4     compute the actual 1-minute stay duration  $S_t$ 
      of EV of type  $v'$  in period  $t$ ;
5     compute  $\bar{x}_{v,t} \leftarrow K/(15\eta) \cdot S_t$ ;
6   end
7   if  $s_v < \tau_v$  then
8     compute the excess charging quantity in arrival
      period  $E \leftarrow \max\{0, x_{v,s_v} - \bar{x}_{v,s_v}\}$ ;
9     update  $x_{v,s_v} \leftarrow x_{v,s_v} - E$ ; set  $t \leftarrow s_v + 1$ ;
10    while  $E > 0$  do
11      update  $x_{v,t} \leftarrow x_{v,t} + E$ ; compute  $E \leftarrow \max$ 
         $\{0, x_{v,t} - \bar{x}_{v,t}\}$ ;
12      update  $x_{v,t} \leftarrow x_{v,t} - E$ ; set  $t \leftarrow t + 1$ ;
13    end
14    compute the excess charging quantity in departure
      period  $E \leftarrow \max\{0, x_{v,\tau_v} - \bar{x}_{v,\tau_v}\}$ ;
15    update  $x_{v,\tau_v} \leftarrow x_{v,\tau_v} - E$ ; set  $t \leftarrow \tau_v - 1$ ;
16    while  $E > 0$  do
17      update  $x_{v,t} \leftarrow x_{v,t} + E$ ; compute  $E \leftarrow \max$ 
         $\{0, x_{v,t} - \bar{x}_{v,t}\}$ ;
18      update  $x_{v,t} \leftarrow x_{v,t} - E$ ; set  $t \leftarrow t - 1$ ;
19    end
20  end
21  for 1-minute period  $t' \in \{s_{v'}, \dots, \tau_{v'}\}$  do
22    set  $t \leftarrow \lceil t'/15 \rceil$ ; set  $y_{v',t'} \leftarrow x_{v,t}/S_t$ ;
23  end
24 end

```

Output: 15minTo1min implementable charging policy $y_{v',t'}$

Specifically, we develop Algorithm 2. In this algorithm, if $s_v = \tau_v$ (i.e., the arrival period is the same as the departure period), the division step (on line 22) is sufficient to convert the 15-minute policy to be 1-minute implementable. Hence, the algorithm needs to shift only when $s_v < \tau_v$ (see line 7 onward). Lines 8–13 (14–19) shift the excess charging quantity from the arrival (departure) period to the middle periods of the charging time window. To demonstrate the effectiveness of Algorithm 2, we compare the following two policies:

- Opt-15minTo1min policy: the 1-minute charging policy transformed via Algorithm 2 from the *optimal* charging policy solved on 15-minute discretization.

- Opt-1-min policy: the *optimal* charging policy solved directly on 1-minute discretization.

Using the optimal policies allows us to isolate the effect of approximation and focus on the performance of Algorithm 2 itself. Therefore, we have to use small 1-minute numerical instances where we can obtain optimal policies by solving SAA with sufficiently many samples. In particular, we choose a one-hour horizon (i.e., 5 a.m. to 6 a.m.), where the number of types with positive arrival rates is 2,439. On such instances, 8,000 samples are sufficient for SAA to generate optimal policies for the SP. Our experiments on other one-hour horizons with a similarly small number of types show the same qualitative results, which are omitted for brevity. For obtaining the Opt-15minTo1min policy, we first compute the optimal 15-minute policy under SAA by generating 8,000 samples from Poisson distributions of customer arrivals with 15-minute discretization and then convert it using Algorithm 2. For obtaining the Opt-1-min policy under SAA, we use 8,000 samples generated from Poisson distributions of customer arrivals with 1-minute discretization. (Our SP model on 15-minute discretization can be easily extended to obtain one on 1-minute discretization.) We then evaluate both the Opt-15minTo1min and Opt-1-min policies on new 10,000 samples generated from the same Poisson distributions with 1-minute discretization. We display the mean and standard deviation of the total cost from both policies in Table 3.

As seen from Table 3, the mean total costs between the Opt-15minTo1min and Opt-1-min policies are close. This shows that our Algorithm 2 is effective in converting a 15-minute policy to a 1-minute implementable policy without losing much optimality.

Last, we show that if we use Algorithm 2 to convert ECP and SAA policies obtained on 15-minute discretization to their respective 1-minute policies, their performance comparison is the same as that under 15-minute discretization. We focus on one-hour horizons where the customer types with positive arrival rates are large, which means the computation is intensive. In particular, we choose 8 a.m. to 9 a.m. and 12 p.m. to 1 p.m., where the number of types is 16,633 and 23,086, respectively.

Table 3. Performance over a One-Hour Time Horizon (5 a.m. to 6 a.m.) with 1-Minute Discretization

Hour	Policy	Total cost		$\mathbb{E}[\text{Demand charge}/\text{total cost}]$
		Mean	Standard deviation	
5 a.m. to 6 a.m.	Opt-1-min	1.046	1.144	0.361
	Opt-15minTo1min	1.061	1.163	0.370

Table 4. Performances of ECP and SAA over Different One-Hour Time Horizons with 1-Minute Discretization

Hour	Policy	Total cost		$\mathbb{E}[\text{Demand charge}/\text{total cost}]$
		Mean	Standard deviation	
8 a.m. to 9 a.m.	ECP-15minTo1min	4.173	2.094	0.340
	SAA(3,000)-15minTo1min	4.191	2.131	0.342
12 p.m. to 1 p.m.	ECP-15minTo1min	19.015	6.938	0.282
	SAA(2,000)-15minTo1min	19.157	7.024	0.287

Note. The number in the parentheses after SAA indicates the number of samples used to compute SAA.

Consequently, SAA cannot be computed with 8,000 samples due to memory limit, and we have to use 3,000 and 2,000 computation samples, respectively. In contrast, ECP can be solved easily as it scales well and does not require samples in computation. (We continue to use 10,000 sample paths for evaluating both policies.) We display the mean and standard deviation of the total expected costs under both ECP-15minTo1min and SAA-15minTo1min in Table 4. As seen, similar to the result on 15-minute discretization in Table 1, in either hour, ECP results in both a lower mean and a lower standard deviation of the total cost than SAA.

In summary, our Algorithm 2 is effective in converting any 15-minute policy to a 1-minute policy, and the performance comparison of these converted policies is consistent with those on 15-minute discretization. Therefore, because of the computational challenges of 1-minute discretization, we focus on 15-minute discretization in our numerical experiments.

7.4.2. Data-Driven Settings. Although in our SP model, we assume Poisson arrival distributions with known arrival rates, we next consider the data-driven setting where there may be estimation errors in the arrival rate (because of various reasons including insufficient data) or mis-specification of arrival distribution (e.g., which may not be truncated Poisson). We consider in such settings whether ECP or SAA is more affected by comparing their out-of-sample performances. We benchmark them with two other approaches. The first approach is empirical risk minimization (ERM), which minimizes the empirical cost on training samples. It differs from SAA used in our paper in that ERM directly uses the training samples to calculate the empirical cost, whereas SAA approximates the expected cost by generating new samples from the truncated Poisson distributions with calibrated arrival rates. The second approach is a robust version of the ECP approach, labeled RECP. Under this approach, each arrival rate is assumed to be in an interval centered at the estimated arrival rate, and then we minimize the worst-case total expected cost, which occurs when the arrival rate is the largest as the cost is nondecreasing in each arrival rate. Specifically, if we denote the ECP model as $\min_{x \in \mathcal{X}} g^E(x, \lambda)$, the RECP model is

obtained by allowing λ to be in the interval $[\underline{\lambda}, \bar{\lambda}]$, where $\underline{\lambda}$ and $\bar{\lambda}$ are the given lower and upper bounds of λ , respectively. Then we have

$$\min_{x \in \mathcal{X}} \max_{\lambda \leq \bar{\lambda}} g^E(x, \lambda) \Leftrightarrow \min_{x \in \mathcal{X}} g^E(x, \bar{\lambda}),$$

where the equivalence of the two formulations is because $g^E(x, \lambda)$ is nondecreasing in λ . We choose $\bar{\lambda} = \hat{\lambda} + M \cdot \left(\frac{\sum_{v \in [V]} \hat{\lambda}_v}{\sqrt{NV}} \right) \mathbf{1}$, where $\hat{\lambda}$ is the estimated arrival rates, M is a hyperparameter, which we set as 0.5 based on cross-validation, and N is the number of training samples.

Estimation Errors in Arrival Rates. We first examine how estimation errors in arrival rates affect the performances of all policies where the arrival distribution is correct. To this end, we use the same distribution (i.e., truncated Poisson) for generating training and testing samples. In particular, we generate $N \in \{50, 100, 200, 365\}$ independent and identically distributed training samples (recall 365 is the number of data points for calibrating the arrival rate of each type in our numerical study), which we use to estimate arrival rates $\hat{\lambda}$. As N increases, there are more data and thus a smaller estimation error. We then feed these estimated arrival rates into ECP, SAA, and RECP to compute their respective policy actions. SAA uses $\hat{\lambda}$ to generate 8,000 new samples to compute its actions, whereas ERM directly uses the N training samples as inputs for the optimization. We then generate from the original truncated Poisson distribution 10,000 testing samples to compute the out-of-sample performances of all policies.

The results are shown in Table 5 for $C = 30$. (Results for when $C = 15, 20$, and 25 are qualitatively similar, which are omitted for brevity.) Among all policies, ERM has the highest total expected cost and the highest associated standard deviation and thus performs the worst. SAA improves significantly on ERM, as different from ERM it can generate new samples to represent scenarios other than the training samples. ECP performs better than both ERM and SAA for any training sample size. This shows that our ECP continues to perform better than SAA (and a related approach ERM) in the presence of estimation errors. The performance of ECP can

Table 5. Policy Performances in the Presence of Estimation Errors When $C = 30$

N	Method	Total cost		$\mathbb{E}[\text{Demand charge}/\text{total cost}]$
		Mean	Standard deviation	
50	ECP	835.39	72.97	0.56
	ERM	876.84	80.20	0.58
	SAA	839.45	73.83	0.56
	RECP	804.88	67.65	0.55
100	ECP	823.28	71.12	0.56
	ERM	875.77	79.49	0.58
	SAA	829.56	72.32	0.56
	RECP	804.14	67.47	0.55
200	ECP	815.55	69.86	0.55
	ERM	867.05	77.93	0.58
	SAA	823.58	71.14	0.56
	RECP	803.81	67.36	0.55
365	ECP	810.29	68.78	0.55
	ERM	857.30	76.10	0.57
	SAA	820.46	70.59	0.56
	RECP	803.58	67.27	0.55

be further improved if we add some robustness to it, that is, using RECP, which significantly improves on ECP, especially when the training sample size N is small (for instance, when $N = 50$ or 100).

To assess how the presence of estimation errors affects the performance of each policy compared with the case without estimation error, we compare Table 5 with the row for $C = 30$ in Table 1. We see that the presence of estimation errors undermines the performances of both ECP and SAA as the total expected costs for

each policy given any N in Table 5 are higher than those in the row for $C = 30$ in Table 1 as expected, especially for small N . However, when $N = 365$ (again the number of data points for calibrating the arrival rate of each type in our numerical study), the mean total costs for ECP and SAA are 810.29 and 820.46, respectively, which are close to those of ECP and SAA with accurate arrival rates in Table 1, that is, 805 and 816.39, respectively. This indicates that though our SP model (where ECP and SAA derive their formulation from) has a large number of customer types and thus may seem over-parameterized, the performance of this model is satisfactory even in the presence of estimation errors.

Model Mis-Specification of Arrival Distribution. We next consider the case when there is model mis-specification, that is, the arrival distribution may not be truncated Poisson. To this end, we use our real arrival data, where the distribution is not necessarily Poisson, and evaluate the performances of all policies to see which one is affected the most.

We repeat the following experiment five times. In each experiment, we randomly equally partition the 365 data points (each corresponding to a day) into five subsamples, where four subsamples are used for training and the remaining one for testing. We use the training subsamples to estimate the arrival rates and feed the estimates into ECP, SAA, and RECP approaches. (Again, ERM directly uses the training subsamples without estimating the arrival rates.) Then we evaluate the solutions obtained for each policy on the testing subsample and report their out-of-sample performances.

Table 6. Policy Performances with Model Mis-Specification When Randomly Partitioning Real Data at $C = 30$

No. of experiment	Method	Total cost		$\mathbb{E}[\text{Demand charge}/\text{total cost}]$
		Mean	Standard deviation	
1	ECP	809.10	102.53	0.55
	ERM	851.41	108.06	0.57
	SAA	818.18	103.24	0.56
	RECP	792.59	95.92	0.54
2	ECP	808.65	101.64	0.55
	ERM	850.69	107.33	0.57
	SAA	818.14	102.36	0.56
	RECP	792.31	94.55	0.54
3	ECP	808.89	101.38	0.55
	ERM	854.26	108.25	0.58
	SAA	815.82	101.21	0.55
	RECP	792.44	95.74	0.54
4	ECP	808.55	101.07	0.55
	ERM	851.60	110.38	0.57
	SAA	816.48	103.44	0.55
	RECP	792.53	95.25	0.54
5	ECP	809.37	102.32	0.55
	ERM	853.07	107.42	0.57
	SAA	816.54	103.11	0.56
	RECP	792.55	96.18	0.54

Table 7. Policy Performances with Model Mis-Specification When Partitioning Real Data in Chronological Order

C	Method	Total cost		$\mathbb{E}[\text{Demand charge}/\text{total cost}]$
		Mean	Standard deviation	
30	ECP	794.89	76.14	0.56
	ERM	831.82	85.90	0.58
	SAA	805.14	72.85	0.57
	RECP	774.37	64.35	0.55
25	ECP	776.61	71.46	0.56
	ERM	819.65	82.42	0.58
	SAA	786.99	72.97	0.56
	RECP	757.07	58.54	0.55
20	ECP	718.94	58.13	0.54
	ERM	763.17	61.79	0.57
	SAA	729.55	54.97	0.55
	RECP	699.37	50.36	0.53
15	ECP	609.64	42.91	0.53
	ERM	657.45	43.06	0.57
	SAA	616.17	42.07	0.54
	RECP	587.95	40.68	0.52

We present the results for each experiment given $C = 30$ in Table 6 (those for other C values are omitted for brevity). As seen, the results are qualitatively very similar to the case with estimation errors in arrival rate. In particular, ERM performs the worst, while SAA improves significantly on ERM. ECP performs better than both ERM and SAA. RECP performs the best. These qualitative results continue to hold if we partition real data into training and testing in chronological order for given different numbers of chargers (see results in Table 7), that is, when we use the first four-fifths of the entire dataset for training and the rest for testing.

In summary, we show that in the presence of arrival rate estimation errors or arrival distribution mis-specification, our ECP approach continues to outperform SAA and a related approach (ERM). The performance of ECP can be significantly improved if we add simple robustness to it. All these results demonstrate the practical use of ECP in data-driven settings.

7.4.3. Adaptive EV Charging Setting. As our ECP policy is static and menu-based that is tailored to customer types, we next examine how well it performs in an adaptive setting where the uncertainty is revealed over time. We benchmark it with a practically implementable adaptive equal charge policy and the rolling horizon implementation of an LP and our ECP.

Equal Charge Policy. This policy is motivated by Blink and ChargePoint. For instance, Blink Charging (2022) “allows equal output to each charger based on the number of stations being used at one time. When one EV is charging, the EV will receive the maximum output.

When others connect, the load will be equally shared between them.” ChargePoint (2022) uses a very similar policy. We implement their “equal charge” (EQC) policy to ensure the feasibility constraints in $\mathcal{X}$ by first determining the total charging rate $X \geq 0$ for all periods and then setting in period t the charging quantity $x_{v,t}$ for each $v \in \mathcal{V}_t$ as

$$x_{v,t} = \min \left\{ K, u_v - \eta \sum_{\ell=s_v}^{t-1} x_{v,\ell}, \max \left\{ \frac{\eta X}{\sum_{v \in \mathcal{V}_t} z_v}, u_v - \eta \sum_{\ell=s_v}^{t-1} x_{v,\ell} - K(\tau_v - t) \right\} \right\} / \eta. \quad (23)$$

(23) shows that $\eta x_{v,t}$ is no larger than any of the three terms: The charging speed limit K (first term), the remaining charging requirement $u_v - \eta \sum_{\ell=s_v}^{t-1} x_{v,\ell}$ (second term), and the third term. The third term itself is the max of two expressions: The first distributes equally the total charging rate ηX to all the EVs available for charging at the station (i.e., $\sum_{v \in \mathcal{V}_t} z_v$); the second, $u_v - \eta \sum_{\ell=s_v}^{t-1} x_{v,\ell} - K(\tau_v - t)$, is the minimum charging quantity in period t to ensure satisfying the charging requirement before the departure time: The difference between the remaining charging requirement (i.e., $u_v - \eta \sum_{\ell=s_v}^{t-1} x_{v,\ell}$) and the maximum possible quantity that can be charged in the remaining periods (i.e., $K(\tau_v - t)$, which is the maximum speed times the number of periods left). When X is large enough, EQC becomes the “maximum-speed charging” policy, which we discuss in Section 7.5.1. Lemma 2 shows that this EQC policy is feasible, that is, $x \in \mathcal{X}$. In the experiment, we adopt a golden-section search procedure to determine the optimal X over the interval $[0, \bar{X}]$, where $\bar{X}$ is a large number.

Lemma 2. Given $X \geq 0$, for x under the EQC policy specified by (23), we have $x \in \mathcal{X}$.

Rolling Horizon Implementation of LP and ECP. We implement the rolling horizon version of a linear program (RH-LP) and our ECP (RH-ECP) similar to Lee et al. (2020). Specifically, in each period, given the system’s status (e.g., the number of EVs at the station and remaining charging requirements), RH-LP resolves the deterministic counterpart of our SP model until the end of the horizon by replacing the future arrivals with their mean values. This deterministic optimization can be rewritten equivalently as an LP. Please find more details and our Algorithm 3 to implement RH-LP in Online Appendix E.1. We implement RH-ECP in a similar rolling fashion and relegate the details and our algorithm to Online Appendix E.2.

We next compare ECP with all three benchmark policies. In particular, we compare ECP with EQC on

Table 8. Performances of ECP and EQC on Large-Scale (96-Period) Charging Instances

C	Method	Total cost		$\mathbb{E}\{\text{Demand charge}/\text{total cost}\}$
		Mean	Standard deviation	
30	ECP	805.00	67.55	0.55
	EQC	810.54	69.93	0.55
25	ECP	783.66	57.13	0.54
	EQC	791.32	60.74	0.54
20	ECP	717.73	44.42	0.53
	EQC	731.87	49.09	0.53
15	ECP	600.40	35.20	0.51
	EQC	620.20	39.59	0.52
10	ECP	441.34	27.71	0.50
	EQC	462.77	31.19	0.52
5	ECP	250.43	19.43	0.51
	EQC	267.75	21.73	0.54

the large-scale (96-period) numerical instances, and with RH-LP and RH-ECP on the small-scale (six-period) instances (discussed in Online Appendix C). This is because even though both RH-LP and RH-ECP can be *computed*

efficiently on large-scale instances and are thus practically implementable, it is computational-prohibitive to *evaluate* both RH-LP and RH-ECP on these large-scale instances. Both are practically implementable because the computation time in each period for RH-LP and RH-ECP, respectively, is 3.75 seconds and 2.8 minutes on average (with 5 minutes at maximum), both of which are sufficient for most practical purposes as they only need to resolve every 15 minutes or longer. However, in terms of evaluation time, recall that we have to evaluate their performances over 10,000 sample paths, and the evaluation on *one* sample path already takes about $3.75 \text{ seconds/period} \times 96 \text{ periods} = 6 \text{ minutes}$ and $2.8 \text{ minutes/period} \times 96 \text{ periods} \approx 4.5 \text{ hours}$, respectively: This is because both approaches need to resolve in each period with the realized uncertainty.

We first show the performances of ECP and EQC evaluated with 10,000 samples in Table 8. Although ECP is static and menu-based, whereas EQC is adaptive based on the realized uncertainty, we find ECP outperforms EQC on each instance. This is because, whereas EQC optimizes only over the total load X , ECP optimizes the entire charging schedule tailored to each

Table 9. Performances of ECP, RH-LP, and RH-ECP on Small-Scale (Six-Period) Instances

C	d	Method	Total cost		$\mathbb{E}\{\text{Demand charge}/\text{total cost}\}$
			Mean	Standard deviation	
12	0.23	ECP	6.00	1.76	0.43
		RH-LP	6.02	1.76	0.46
		RH-ECP	5.99	1.72	0.43
	0.46	ECP	8.64	2.53	0.60
		RH-LP	8.81	2.63	0.63
		RH-ECP	8.59	2.46	0.60
	0.93	ECP	13.91	4.08	0.75
		RH-LP	14.40	4.40	0.77
		RH-ECP	13.77	3.93	0.75
10	0.23	ECP	5.75	1.52	0.44
		RH-LP	5.77	1.53	0.46
		RH-ECP	5.71	1.48	0.43
	0.46	ECP	8.28	2.18	0.60
		RH-LP	8.44	2.29	0.63
		RH-ECP	8.18	2.10	0.60
	0.93	ECP	13.49	3.56	0.75
		RH-LP	13.80	3.85	0.77
		RH-ECP	13.12	3.37	0.75
8	0.23	ECP	5.41	1.25	0.43
		RH-LP	5.26	1.22	0.47
		RH-ECP	5.23	1.19	0.43
	0.46	ECP	8.08	1.84	0.61
		RH-LP	7.72	1.83	0.64
		RH-ECP	7.59	1.69	0.60
	0.93	ECP	13.00	2.95	0.76
		RH-LP	12.65	3.09	0.78
		RH-ECP	12.15	2.70	0.75

Table 10. Objective Values and Performances of ECP-C and No-Ent

C	Method	Objective	Total cost	
			Mean	Standard deviation
30	ECP-C	899.21	805.00	67.55
	No-Ent	1,366.00	834.90	71.43
25	ECP-C	891.43	783.66	57.13
	No-Ent	1,227.68	818.03	60.40
20	ECP-C	854.78	717.73	44.42
	No-Ent	1,082.74	756.78	46.25
15	ECP-C	779.44	600.40	35.20
	No-Ent	916.34	639.09	36.35
10	ECP-C	639.18	441.34	27.71
	No-Ent	701.80	478.87	29.32
5	ECP-C	411.01	250.43	19.43
	No-Ent	421.15	282.86	19.72

customer type. The superiority of ECP over EQC is much more pronounced in long-duration charging instances, which we omit for brevity.

We next compare ECP with RH-LP and RH-ECP on the small-scale (six-period) numerical instances in Table 9. We find that RH-ECP performs the best among all the policies in terms of both the mean and standard deviation of the total cost. When C is large, even the (static) ECP performs better than the sensible adaptive policy RH-LP on many instances. This suggests our proposed ECP approach is sufficiently adaptive to customer types (arrival/departure time and charging requirement) and time period.

7.4.4. Role of Entropic Dominance Constraints. We now examine the role of entropic dominance constraints in constructing our ECP formulation. Recall in Section 4.2, we obtain our ECP by leveraging the DRO framework and use the ambiguity set $\mathcal{F}$, which is the intersection of $\mathcal{F}^1$, containing the entropic dominance constraints, and $\mathcal{F}^2$, a commonly used ambiguity set with the support and mean information. To examine the role of $\mathcal{F}^1$, we compare the DRO formulation with $\mathcal{F}^1$, which is DRO-Ent (in turn approximated by ECP-C as seen in Section 4.2), with another DRO formulation without $\mathcal{F}^1$ as follows:

$$\min_{x \in \mathcal{X}} \sup_{\mathbb{P} \in \mathcal{F}^2} \mathbb{E}_{\mathbb{P}}[c(x, \tilde{z})]. \quad (\text{No-Ent})$$

In other words, No-Ent is the same as DRO-Ent except that the ambiguity set is $\mathcal{F}^2$ instead of $\mathcal{F}$. Hence, we need to compare ECP-C with No-Ent, which can be reformulated as a tractable LP.

Table 10 compares the objective values and performance evaluations of ECP-C and No-Ent given different values of the number of chargers C . We see that as C decreases, the objective value of ECP-C tends to coincide

Table 11. Performance Improvement of ECP-C over the Maximum-Speed Charging Policy

C	Total cost		Expected maximum load			
	Mean	Standard deviation	Overall	On-peak	Mid-peak	Off-peak
15	13.60%	27.87%	23.64%	23.56%	21.43%	10.33%
20	12.12%	25.49%	20.99%	20.68%	18.56%	8.41%
25	11.22%	21.80%	18.91%	19.00%	16.03%	8.70%
30	10.99%	19.02%	17.94%	18.89%	13.65%	8.59%


Notes. The second and third columns are the percentage reduction in the mean and standard deviation, respectively, of the total cost of ECP-C compared with the maximum-speed charging policy. Analogously, the fourth to seventh columns are the percentage difference of the mean of the maximum load over the entire time horizon, on-peak hours, mid-peak hours, and off-peak hours, respectively. The expected maximum load over the entire horizon (fourth column) differs from the maximum of the expected maximum load during different time intervals (in the fifth to seventh columns), as the maximum load may occur in different time intervals.

with that of No-Ent. For instance, when $C=5$, their objective values are very close, differing by about 2%. Nevertheless, when the solutions from ECP-C and No-Ent are evaluated on 10,000 sample paths, the mean total cost of ECP-C is about 11% smaller than that of No-Ent. This observation holds for all values of C considered in the table. This highlights the important role entropic dominance constraints play in the ambiguity set in constructing ECP.

7.5. Managerial Insights in EV Charging Scheduling

We use ECP-C to first study the value of customer desired departure time information and then examine the effect of the composition of demand charge on electricity load.

7.5.1. Value of Desired Departure Time Information. We study this value by comparing the performance of ECP-C with that of the maximum-speed charging policy, which charges vehicles using maximal charging speed (power capacity) until completion on a first-come-first-served basis and is often used in practice and literature (Zhang et al. 2021). Denote by x^M the charging decision under this maximum-speed charging policy. Then $x_{v,t}^M = \max\{0, \min\{K, u_v - (t - s_v) \cdot K\}\}$ for all $v \in [V]$ and $t \in \mathcal{T}_v$. We evaluate the maximum-speed charging policy in terms of both the mean and standard deviation of the total cost analogous to the evaluation of ECP-C. We also evaluate under both policies the mean of the maximum electricity load during the entire horizon, on-peak hours, mid-peak hours, and off-peak hours. Table 11 lists all the performance measures of ECP-C relative to those of the maximum-speed charging policy. Using column 2 as an example, if we denote $\pi^M = \mathbb{E}_{\mathbb{P}^C} [c(x^M, \tilde{z})]$, column 2 represents $(\pi^M - \pi^E)/\pi^M \times 100\%$.

Figure 5. (Color Online) Electricity Load Under ECP-C for Different Compositions of the Demand Charge Given $C = 30$ 

Note. The black, white, and gray regions on the x axis represent on-peak hours (hours 13 to 18), mid-peak hours (hours 9 to 12 and hours 19 to 23), and off-peak hours (hours 1 to 8 and hour 24), respectively.


As seen in Table 11, compared with the maximum-speed charging policy, ECP-C results in a lower mean and standard deviation of the total cost, for example, with 30 chargers, by 10.99% and 19.02%, respectively. This lower cost of ECP-C is reflected in the lower expected maximum electricity load during any time interval, implying a more smooth load. Moreover, the lower expected maximum load from ECP-C alleviates the potential overcharge to EVs under the maximum-speed charging policy, and thus ECP-C can also lower the risk of EV battery degradation (Thompson 2018). When station capacity decreases, the difference between ECP-C and the maximum-speed charging policy increases. This difference can be regarded as the value of information on the customer departure time, as ECP-C differs from the maximum-speed charging policy in that whereas the former uses this information, the latter does not. This value is even higher for the long-duration charging instances, which we omit for brevity. These results suggest that a charging station can benefit significantly from incorporating the customer desired departure information, as our SP model does, to smooth the electricity load and reduce the total cost.

7.5.2. Effect of the Composition of the Total Demand Charge on Electricity Load. The rate structure of most utility firms has three common components of time-of-

use (TOU) demand charges, namely on-peak, mid-peak, and all-period demand charges (see their corresponding hours in Figure 5). We examine the following compositions of the total demand charge: (i) without any demand charge; (ii) with only on-peak demand charge; (iii) with only mid-peak demand charge; (iv) with both on-peak and mid-peak demand charges; (v) with only all-period demand charge; and (vi) with both all-period and on-peak demand charges. We compare these with the baseline case where all three components exist. Given each composition of the total demand charge, we compute x^E , the optimal solution to ECP-C, and then use it to compute the mean and standard deviation of the electricity load in the charging station for each period, which we plot in Figure 5. We make the following observations.

Observation 1. In the absence of any demand charge, the TOU energy charge creates an artificial electricity load spike around the transitions between any two time intervals. As seen in Figure 5(a), the expected load swings significantly around the hours between any two sets of time intervals: around hour 9 (between off-peak and mid-peak hours) and around hours 12 and 18 (between mid-peak and on-peak hours). This means that the TOU energy charge (see tariff structure in Section 7.1)

Figure 6. (Color Online) Performance of the Alternating Optimization Algorithm on JPS-C



Notes. The objective values on the dotted line at integer number of iterations correspond to the optimal values of the scheduling optimization, whereas those at the half-integer number of iterations correspond to the optimal values of the pricing optimization. The asterisks (*) correspond to the expected profit in (17) evaluated at JPS-C solutions after each iteration.

alone is not sufficient to smooth the electricity load and may even create artificial spikes, putting undue stress on the electric grid. This is consistent with the observation in a similar setting in practice (Zhang and Qian 2018): At midnight, there is a huge spike in electricity charging load from residential customers, because their electricity tariff structures do not include any demand charge while their electricity price at midnight is lower than that at hour 11 p.m. We show this observation holds when the decisions are near optimal in the EV charging management setting.

Observation 2. *Either an on-peak or mid-peak demand charge alone cannot smooth the electricity load and may even create a larger load swing than the case without any demand charge.* As shown in Figure 5(b), when only an on-peak demand charge exists, the electricity load is smoothed only within on-peak hours but not over time. In particular, there are some artificial spikes in the load around all transition hours (i.e., between on-peak and mid-peak hours and between mid-peak and off-peak hours) because the charging station simply shifts the high load from on-peak hours to other hours. These spikes are even higher than those within the same hours when there is no demand charge, and thus impose more stress on the electric grid. When there is only a mid-peak demand charge, an analogous pattern occurs (the corresponding figure is omitted for brevity). Similarly, when both an on-peak demand charge and mid-peak demand charge exist, the load is smoothed only during these combined hours, and load spikes occur around hours between mid-peak and off-peak hours (figure omitted for brevity).

Observation 3. *An all-period demand charge alone is sufficient to smooth electricity load over time.* In contrast to the case in which there is no demand charge or only either an on-peak or mid-peak demand charge, an all-period


demand charge results in a smooth electricity load (Figure 5(c)). The load curve for this case almost coincides with that for the case in which all three components of demand charge exist (Figure 5(d)). This is because, with an all-period demand charge alone, the charging station is incentivized to reduce the highest load over the entire time horizon.

7.6. Joint Pricing and Scheduling

We use both ECP and SAA approaches to solve the joint pricing and scheduling problem in (17). For simplicity, we set $C = 30$; we use only all-period demand charge and set $d = \$0.930/\text{kWh}$; we set $[p, \bar{p}]$ as $[0.3, 0.7]$, $r_v = 4/3$ for all $v \in [V]$, and $\bar{\lambda}_v$ so that the resulting arrival rates when $p = 0.5$ (the midpoint of $[p, \bar{p}]$) are the same as those in Section 7.1. We set the stopping tolerance $\delta = 0.01$.

We solve (17) using JPS-C with an alternating optimization algorithm analogous to Algorithm 1. Figure 6 plots the objective values of JPS-C and evaluations of JPS-C (i.e., substituting the JPS-C solutions to (17) on 10,000 sample paths) at two different initial values of p : 0.3 and 0.7. We find that the algorithm converges in four iterations to give the same price $p = 0.49$. Hence, our method is fast and not sensitive to the initial price p . Although our algorithm is not guaranteed to converge to a globally optimal solution of JPS-C, it performs well on our numerical instances, where $p = 0.49$ can be numerically verified as the globally optimal price.

We also solve (17) using SAA, for which we cannot directly extend (14) as the arrival rates of Poisson random variables, λ , depend on the pricing decision, p . Thus, we first enumerate p and solve for each value of p the corresponding SAA problem to obtain an expected profit, and finally select the value of p that results in the highest expected profits. We use 8,000 samples and enumerate p from 0.3 to 0.7 with a stepsize of 0.005. Figure 7 plots the

Figure 7. (Color Online) Expected Profits Under SAA and JPS-C

expected profits under SAA at different $p \in [0.45, 0.54]$ (while the rest of the range is omitted), which shows that the optimal p under SAA is 0.48. (JPS-C does not need to enumerate p and instead uses the alternating algorithm mentioned previously, so there is only one objective value under JPS-C in Figure 7.)

We evaluate SAA by substituting its solutions into (17) on 10,000 samples to obtain the corresponding expected profits. The optimal price $p = 0.48$ under SAA leads to an evaluation of the expected profits under SAA to be 739.62, lower than the expected profits under JPS-C, which is 749.91. The difference between these two is significant as the standard error of the expected profits under both is 0.72, which is very low compared with the difference between their expected profits. The superiority of JPS-C is due to ECP-C's better performance in scheduling optimization than SAA, as mentioned in Section 7.3. Moreover, JPS-C runs much faster, whereas SAA requires solving 81 large-scale SAA instances, each of which can be time-consuming. In summary, ECP (or JPS-C) is a better approach than SAA in solving the joint pricing and scheduling problem.

8. Conclusions and Future Work

We use an ECP approach to examine the EV charging management of a service provider who faces uncertain arrivals of customers with heterogeneous arrival/desired departure times and charging requirements. We formulate this problem as an SP where customer arrivals follow a joint Poisson distribution to minimize the expected total electricity cost that includes demand charges. As the SP is large-scale, we develop ECP approximations. When the service provider is uncapacitated, we develop an ECP approximation by bounding the largest order statistic, introduced by the demand charge, using exponential cones based on the epigraph of MGFs of the Poisson customer arrivals. We show that this ECP gives an upper bound of the SP and characterize a bound on the gap

between their theoretical performances. In the capacitated case, we develop another ECP approximation by also using the idea from DRO of employing an entropic dominance ambiguity set. Rather than mitigating distributional ambiguity, we use DRO to derive an ECP as a tractable upper bound of the corresponding SP. Moreover, we extend our ECP approach to jointly optimize pricing and scheduling and propose an alternating optimization algorithm to solve the resulting ECP approximations efficiently.

We then benchmark our ECP approach with two common approaches in solving SPs, SAA and DRO, on numerical instances calibrated to real tariff and EV customer arrival data. As our numerical instances are large-scale (with about 80,000 random variables and 700,000 decision variables), we show that DRO cannot be solved on these instances; ECP runs about 50 times faster than SAA and results in not only a lower standard deviation of the total cost but also a lower mean of this cost because SAA does not scale well and cannot be solved to optimality in a reasonable amount of time. We also find the optimality gap of ECP is at most around 4% on all instances. The superior performance of ECP over SAA continues to hold when considering three practical implementation issues: a finer time discretization, a data-driven setting when there are estimation errors in the arrival rate (because of insufficient data) or even the mis-specification of the arrival distribution, and an adaptive charging setting where uncertainty is revealed over time. In the extended model of joint pricing and scheduling, ECP results in a higher expected profit than SAA and also runs faster. In addition, ECP is a deterministic approximation, whereas SAA depends on samples generated. Therefore, ECP emerges as the best approach in solving the EV charging problem.

We finally use our ECP approach to generate insights on a similar set of data-calibrated numerical instances and find the following: First, it is beneficial for the charging service provider to use the customer departure information, as the ECP approach outperforms by about 11% a commonly used maximum-speed charging policy that ignores such information. Second, a TOU energy charge alone (i.e., without any demand charge) creates artificial spikes in the electricity load, consistent with the observation in practice, which we show to hold when the decisions are near optimal in the EV management setting. Whereas either an on-peak or mid-peak demand charge shifts the artificial spikes and could not smooth the electricity load, an all-period demand charge alone can result in a load as smooth over time as the case in which all three types of demand charge exist.

Our ECP approach is appealing for use in practice by a charging service provider as it outperforms various benchmarks considering different practical implementation issues. The managerial insights can give guidance for charging service providers in terms of the value of

collecting departure time information from customers and policymakers in designing the electricity tariff structure to smooth electricity load.

There are a few promising avenues to extend this work. First, our SP model assumes that the departure time provided by customers is accurate. In practice, customers may underreport their departure time to reduce the risk of insufficient charge upon pickup, and thus it is potentially beneficial for the service provider to charge less aggressively to minimize the total expected cost. Considering this behavior should prove to be fruitful. Second, our methods to construct ECP approximations may be applicable to other two-stage stochastic linear programs with fixed recourse.

Acknowledgements

The authors thank the entire review team for the wonderful suggestions that helped significantly improve the paper.

Endnotes

¹ See <https://ev.caltech.edu/info>.

² For customers arriving at the same time, we allow any way of truncation.

³ All of our analytical results can be extended in a straightforward manner when we modify our model to represent a steady-state situation, where type v with $\tau_v < s_v$ represents customers departing in period $T + \tau_v$.

⁴ For the case in which η and K differ for different customer types (e.g., with various states of charge), our theoretical results in Sections 4 and 6 extend in a straightforward manner.

⁵ We also tried time discretization by rounding up the stay duration to the integer multiple of 15 minutes, that is, $\lceil \text{stay duration}/15 \rceil \times 15$. We find that it performs worse as it gives a higher total expected cost for both ECP and SAA.

⁶ Details are omitted for brevity and are available upon request.

References

- Avci B, Girotra K, Netessine S (2014) Electric vehicles with a battery switching station: Adoption and environmental impact. *Management Sci.* 61(4):772–794.
- Ben-Tal A, El Ghaoui L, Nemirovski A (2009) *Robust Optimization*, vol. 28 (Princeton University Press, Princeton, NJ).
- Ben-Tal A, Den Hertog D, De Waegenaere A, Melenberg B, Rennen G (2013) Robust solutions of optimization problems affected by uncertain probabilities. *Management Sci.* 59(2):341–357.
- Bertsimas D, Sim M (2004) The price of robustness. *Oper. Res.* 52(1):35–53.
- Bertsimas D, Sim M, Zhang M (2019) Adaptive distributionally robust optimization. *Management Sci.* 65(2):604–618.
- Birge JR, Louveaux F (2011) *Introduction to Stochastic Programming* (Springer Science & Business Media, Boston).
- Blink (2020) EV charging rates. Accessed December 2020, <https://www.blinkcharging.com/drivers/pricing/>.
- Blink Charging (2022) Blink and local load management. Accessed March 2022, <https://blinkcharging.com/blink-and-local-load-management/>.
- Bloomberg (2018) U.S. electric vehicle charging market to grow to \$18.6 billion. Accessed December 2019, <https://www.bloomberg.com/news/articles/2018-08-02/u-s-electric-vehicle-charging-market-to-grow-to-18-6-billion>.
- Boyd S, Vandenberghe L (2004) *Convex Optimization* (Cambridge University Press, Cambridge, UK).
- Chandrasekaran V, Shah P (2017) Relative entropy optimization and its applications. *Math. Programming* 161(1–2):1–32.
- Chares R (2009) Cones and interior-point algorithms for structured convex optimization involving powers and exponentials. PhD thesis, UCL-Université Catholique de Louvain, Louvain-la-Neuve, Belgium.
- ChargePoint (2022) Charge more EVs with power management. Accessed March 2022, <https://www.chargepoint.com/blog/charge-more-evs-power-management>.
- Chen Z, Sim M, Xu H (2019) Distributionally robust optimization with infinitely constrained ambiguity sets. *Oper. Res.* 67(5):1328–1344.
- Chitkara A, Cross-Call D, Li B, Sherwood J (2016) A review of alternative rate designs. Accessed December 2019, <https://rmi.org/wp-content/uploads/2017/04/A-Review-of-Alternative-Rate-Designs-2016.pdf>.
- Choi S, Ruszczyński A (2008) A risk-averse newsvendor with law invariant coherent measures of risk. *Oper. Res. Lett.* 36(1):77–82.
- Danzig GB (1955) Linear programming under uncertainty. *Management Sci.* 1(3–4):197–206.
- Delage E, Ye Y (2010) Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Oper. Res.* 58(3):595–612.
- Esfahani PM, Kuhn D (2018) Data-driven distributionally robust optimization using the wasserstein metric: Performance guarantees and tractable reformulations. *Math. Programming* 171(1–2):115–166.
- Glassmire J, Komor P, Lilienthal P (2012) Electricity demand savings from distributed solar photovoltaics. *Energy Policy* 51:323–331.
- Goldfarb D, Iyengar G (2003) Robust portfolio selection problems. *Math. Oper. Res.* 28(1):1–38.
- Hanasusanto GA, Kuhn D, Wiesemann W (2016) A comment on “computational complexity of stochastic programming problems.” *Math. Programming* 159(1–2):557–569.
- Hao Z, He L, Hu Z, Jiang J (2020) Robust vehicle pre-allocation with uncertain covariates. *Production Oper. Management* 29(4):955–972.
- Hausman WJ, Neufeld JL (1984) Time-of-day pricing in the us electric power industry at the turn of the century. *RAND J. Econom.* 15(1):116–126.
- He L, Ma G, Qi W, Wang X (2021) Charging an electric vehicle-sharing fleet. *Manufacturing Service Oper. Management* 23(2):471–487.
- Jaillet P, Loke GG, Sim M (2022) Strategic workforce planning under uncertainty. *Oper. Res.* 70(2):1042–1065.
- Jiang DR, Powell WB (2016) Practicality of nested risk measures for dynamic electric vehicle charging. Preprint, submitted May 10, <https://arxiv.org/abs/1605.02848>.
- Jin C, Tang J, Ghosh P (2013) Optimizing electric vehicle charging with energy storage in the electricity market. *IEEE Trans. Smart Grid* 4(1):311–320.
- J. P. Morgan (2018) Driving into 2025: The future of electric vehicles. Accessed December 2019, <https://www.jpmorgan.com/global/research/electric-vehicles>.
- Karlsson S, Kushnir D (2013) How energy efficient is electrified transport. Sandén B, ed. *Systems Perspectives on Electromobility* (Chalmers University of Technology, Gothenburg, Sweden), 45–59.
- Lee ZJ, Pang JZ, Low SH (2020) Pricing EV charging service with demand charge. *Electric Power Systems Res.* 189:106694.
- Lim MK, Mak HY, Rong Y (2014) Toward mass adoption of electric vehicles: Impact of the range and resale anxieties. *Manufacturing Service Oper. Management* 17(1):101–119.
- Mak HY, Rong Y, Shen ZJM (2013) Infrastructure planning for electric vehicles with battery swapping. *Management Sci.* 59(7):1557–1575.
- McKinsey & Company (2018) Charging ahead: Electric-vehicle infrastructure demand. Accessed January 2020, <https://perma.cc/9P7A-SP8Q>.
- Mendelson H, Whang S (1990) Optimal incentive-compatible priority pricing for the M/M/1 queue. *Oper. Res.* 38(5):870–883.
- Mills A, Wiser R, Barbose G, Golove W (2008) The impact of retail rate structures on the economics of commercial photovoltaic systems in California. *Energy Policy* 36(9):3266–3277.

- MOSEK ApS (2020) MOSEK Fusion API for Python 9.2. Accessed January 2021, <https://docs.mosek.com/9.2/pythonfusion/index.html>.
- National Renewable Energy Laboratory (2017) Identifying potential markets for behind-the-meter battery energy storage: A survey of U.S. demand charges. Accessed January 2020, <https://www.nrel.gov/docs/fy17osti/68963.pdf>.
- Nesterov Y, Nemirovskii A (1994) *Interior-Point Polynomial Algorithms in Convex Programming* (SIAM, Philadelphia).
- Neubauer J, Simpson M (2015) Deployment of behind-the-meter energy storage for demand charge reduction. Accessed December 2019, <https://www.nrel.gov/docs/fy15osti/63162.pdf>.
- Nissan (2019) 2019 Nissan LEAF range, charging & battery. Accessed December 2019, <https://www.nissanusa.com/vehicles/electric-cars/leaf/features/range-charging-battery.html>.
- Powerflex (2021) Powerflex product guide 2021. Turnkey electric vehicle charging solutions. Accessed May 2022, <https://www.powerflex.com/wp-content/uploads/PF-Product-Guide-2021-FINAL-Digital.pdf>.
- Schneider F, Thonemann UW, Klabjan D (2017) Optimization of battery charging and purchasing at electric vehicle battery swap stations. *Transportation Sci.* 52(5):1211–1234.
- See CT, Sim M (2010) Robust approximation to multiperiod inventory management. *Oper. Res.* 58(3):583–594.
- Shapiro A, Nemirovski A (2005) On complexity of stochastic programming problems. Vaithilingam J, Rubinov A, eds. *Continuous Optimization: Current Trends and Modern Applications* (Springer, Boston), 111–146.
- Shapiro A, Dentcheva D, Ruszczyński A (2009) *Lectures on Stochastic Programming: Modeling and Theory* (SIAM, Philadelphia).
- Sun B, Sun X, Tsang DH, Whitt W (2019) Optimal battery purchasing and charging strategy at electric vehicle battery swap stations. *Eur. J. Oper. Res.* 279(2):524–539.
- Tesla (2020) Tesla supercharging. Accessed December 2020, <https://www.tesla.com/support/supercharging/>.
- Thompson AW (2018) Economic implications of lithium ion battery degradation for Vehicle-to-Grid (V2X) services. *J. Power Sources* 396:691–709.
- UK Department for Transport (2018) Electric chargepoint analysis 2017: Local authority rapids (revised). Accessed December 2019, <https://www.gov.uk/government/statistics/electric-chargepoint-analysis-2017-local-authority-rapids>.
- Veall MR (1983) Industrial electricity demand and the Hopkinson rate: An application of the extreme value distribution. *Bell J. Econom.* 14(2):427–440.
- Wiesemann W, Kuhn D, Sim M (2014) Distributionally robust convex optimization. *Oper. Res.* 62(6):1358–1376.
- Wood E, Rames C, Muratori M, Raghavan S, Melaina M (2017) National plug-in electric vehicle infrastructure analysis. Accessed January 2020, <https://www.nrel.gov/docs/fy17osti/69031.pdf>.
- Wu OQ, Yücel Ş, Zhou Y (2022) Smart charging of electric vehicles: An innovative business model for utility firms. *Manufacturing Service Oper. Management* 24(5):2481–2499.
- Xu Y, Pan F, Tong L (2016) Dynamic scheduling for charging electric vehicles: A priority rule. *IEEE Trans. Automated Control* 61(12):4094–4099.
- Zap-Map (2019) EV connector types. Accessed December 2019, <https://www.zap-map.com/charge-points/connectors-speeds/>.
- Zhang G, Tan ST, Wang GG (2017) Real-time smart charging of electric vehicles for demand charge reduction at non-residential sites. *IEEE Trans. Smart Grid* 9(5):4027–4037.
- Zhang P, Qian ZS (2018) User-centric interdependent urban systems: Using time-of-day electricity usage data to predict morning roadway congestion. *Transportation Res. Part C Emerging Tech.* 92:392–411.
- Zhang T, Chen W, Han Z, Cao Z (2014) Charging scheduling of electric vehicles with local renewable energy under uncertain electric vehicle arrival and grid power price. *IEEE Trans. Vehicular Tech.* 63(6):2600–2612.
- Zhang Y, Augenbroe G (2018) Optimal demand charge reduction for commercial buildings through a combination of efficiency and flexibility measures. *Appl. Energy* 221:180–194.
- Zhang Y, Lu M, Shen S (2021) On the values of vehicle-to-grid electricity selling in electric vehicle sharing. *Manufacturing Service Oper. Management* 23(2):488–507.
- Zhu T, Xie J, Sim M (2021) Joint estimation and robustness optimization. *Management Sci.* 68(3):1659–1677.

Li Chen is a research fellow in the Institute of Operations Research and Analytics at National University of Singapore. He has broad interests in prescriptive analytics, particularly in decision making and optimization under uncertainty with applications in operations.

Long He is an associate professor of decision sciences at the School of Business, George Washington University. His current research involves using data-driven approaches to address problems in smart city operations (e.g., electric vehicles, vehicle sharing, last-mile delivery) and supply chain management.

Yangfang (Helen) Zhou is an associate professor of operations management at the Lee Kong Chian School of Business, Singapore Management University. Her main research interest is sustainable operations, focusing on renewable energy, electric mobility, sustainable agriculture, and food waste.