

Infrastructure Planning for Electric Vehicles with Battery Swapping

Ho-Yin Mak

Department of Industrial Engineering and Logistics Management, The Hong Kong University of Science and Technology,
Kowloon, Hong Kong
hymak@ust.hk

Ying Rong

Antai College of Economics and Management, Shanghai Jiao Tong University, Shanghai 200052, China
yrong@sjtu.edu.cn

Zuo-Jun Max Shen

Department of Industrial Engineering and Operations Research, University of California at Berkeley, Berkeley CA 94720,
shen@ieor.berkeley.edu

The transportation sector is a major source of greenhouse gas (GHG) emissions. As a step toward a greener environment, solutions involving electric vehicles (EVs) have been proposed and discussed. When powered by electricity from efficient and environmentally-friendly generators, EVs have significantly lower per-mile running costs compared to gasoline cars, while generating lower emissions. Unfortunately, due to the limited capacity of batteries, typical EVs can only travel for about 100 miles on a single charge. Because recharging takes several hours, it is impossible to recharge an EV in the middle of a long (round) trip exceeding 100 miles. Better Place (BP), a start-up based in Palo Alto, CA, proposed a novel strategy that potentially overcomes the recharging problem. In the plan, in addition to charging adaptors at homes, work places and shopping malls, “swapping stations”, at which depleted batteries can be exchanged for recharged ones in the middle of long trips, will be located at strategic locations along freeways. With its battery swapping equipment, BP has demonstrated how to effectively refuel an EV in less than two minutes. The possible success of EV solutions based on the idea of battery swapping hinges on the ability of the charging service provider (BP or other similar firms) to deploy a cost-effective infrastructure network with comprehensive coverage. Unfortunately, since the adoption rate of electric vehicles, and thus demand for swapping service, is still highly uncertain, the service provider must make deployment plans with incomplete information on hand. In this paper, we develop models that aid the planning process for deploying battery swapping infrastructure, based on a robust optimization framework. We further show that our models can be tightly approximated by mixed-integer second-order cone programs (MISOCPs), which are readily solvable by commercial solvers. Using these models, we demonstrate the potential impacts of battery standardization and various technology advancements on the optimal infrastructure deployment strategy.

Key words: Electric Vehicles, Green Transportation, Infrastructure Investment, Robust Optimization, Facility Location

1. Introduction

The transportation sector holds the key to a greener environment. In 2003, the sector was accountable for 29.6% of the total greenhouse gas (GHG) emissions in the US. GHG emissions by transportation also grew the fastest among all economic sectors, accounting for 47% of the net increase in total emissions since 1990 [50]. Such high emissions by the sector can be largely attributed to the heavy dependence on fossil fuels. To cut down emissions by land transportation, the industry has come up with solutions involving alternative fuel vehicles (AFVs). In particular, plug-in electric vehicles (EVs) provide the possibility to completely switch away from the dependence on oil.

Compared to internal combustion engine (ICE) vehicles, EVs deliver a number of potential benefits. The first and most obvious is the lower emissions. EVs have no tailpipe emissions and thus produce less pollution on the end-user side. Even considering the total emissions in the entire electricity supply chain, because of the usage of cleaner and more efficient power generators, especially in more developed countries, the total emissions by EVs are still significantly lower than by ICE vehicles. For example, an EV that is powered by the current US electric grid causes 50% less CO₂ emissions per mile travelled, well-to-wheel (i.e., including the fuel supply chain footprint), compared to an ICE vehicle [37]. Second, EVs may be powered by renewable energy sources. Therefore, EV operations are relatively insensitive to factors such as the depletion of fossil fuels and supply uncertainty of crude oil. Furthermore, on the consumer side, the cost of energy in the form of electricity is cheaper than in the gasoline form. In the US, the end-user costs are roughly 2-3 cents per mile for EVs and 13 cents per mile for ICE vehicles [30]. Even including the additional storage cost needed to purchase a battery, which is amortized to 6 cents per mile [26], the total energy usage cost for EVs is still significantly lower.

However, there are several major challenges to overcome before EVs will be adopted on a mass scale. Perhaps the most significant obstacle comes from the nature of recharging. While ICE cars can be easily refueled at accessible gas stations in a matter of minutes, EVs require specific equipment and a significant amount of time (in the order of hours) to be recharged. The shorter range of EVs than that of ICE cars, which necessitates more frequent recharges than gasoline refuels, only makes this problem worse. While it is not difficult to park the EVs at work place or home for a few hours or over night, it is impractical to station the car for a recharge in the middle of a long trip.

Better Place (BP), a venture-backed company based in Palo Alto, CA, proposed an innovative infrastructure strategy that potentially overcomes the above obstacles. In the plan, in addition to charging adaptors at homes, work places and shopping malls, the network of support facilities for EVs will also include swapping stations at which EVs may exchange their depleted batteries for full ones in the middle of long trips. The depleted batteries are then recharged at the stations,

and later swapped for depleted batteries on other arriving EVs. This solution allows EVs to be effectively “refueled” in only 1-2 minutes instead of hours.

The battery swapping concept is an integral component of BP’s novel business model, in which its role is somewhat analogous to cell phone service providers. It builds and operates the infrastructure system (i.e., charging adaptors and battery swapping stations) and provides service (“refueling” EVs). Because batteries can be swapped, they are not owned by the EV users themselves, but rather leased to them based on some service contracts. Users will be charged based on usage (i.e., miles driven), much like cell phone users being billed on talk time. This arrangement poses several potential benefits. It decouples ownerships of the battery and the EV, two items with different life cycles, and makes it very easy to take advantage of future improvements in battery technology by regular battery replacements, compared to other solutions that bundle the EV and battery as one single unit. Besides, the significantly lower one-time cost for the user to purchase a new EV (the battery typically costs around \$10,000) can help boost adoption rates. The innovative strategic plan of BP has attracted much attention and venture capital [45]. Under favorable policies by respective governments, the company is beginning to implement their projects in Israel and Denmark, and is planning to extend coverage to Australia, Canada, China, Japan, and the US (California and Hawaii).

The infrastructure deployment process for EVs involves a chicken-and-egg dilemma. On one hand, drivers are reluctant to purchase EVs until the coverage of recharging infrastructure is comprehensive enough. On the other hand, charging service providers are not willing to invest heavily on infrastructure unless significant demand is observed. To address this issue, BP attempts to start rolling out its network of recharging facilities before the vehicles are scheduled to be delivered [12]. Therefore, BP must make the strategic investment decisions of where and how to deploy the infrastructure well before observing the real subscription numbers.

BP’s infrastructure system will consist of two main types of facilities: charging adaptors, at which users park their EVs for a few hours to recharge their batteries, and swapping stations, at which users swap batteries with minimal delay on their trips. The locations of charging adaptors are relatively easy to determine, because users are likely to park their cars for hours only at selected locations such as their homes, work places and major shopping centers. Naturally, regular usage, such as shopping or commuting to school or work, can be mostly covered by these charging adaptors. Swapping stations, on the other hand, are needed for long trips (or round-trips) exceeding the range of a full battery. To ensure full coverage of the freeway network, swapping stations must be located at convenient locations along such long trips. When deciding where to locate swapping stations, BP must also keep in mind that the demand rate for swapping at a station determines the number of spare batteries that the station must hold to ensure that users can almost always obtain

full batteries without waiting. Because batteries are expensive and have limited life, holding costs are significant. Furthermore, having a higher demand rate and more spare batteries at a single station means that more batteries will be recharged in parallel and the load on the electric grid will be heavier. BP must take into account all these factors when designing the network of swapping stations.

Unfortunately, due to the need to deploy a complete infrastructure network before delivering the vehicles, BP has to make strategic network design decisions before observing the actual demand. At this point, factors affecting swapping demand, such as the EV adoption rate, BP's market share and the driving patterns of EV users, are not precisely known. With only limited information, such as the sizes of the commercial fleets of corporate partners having signed contracts with it [12], and demand forecasts, BP has to come up with a robust swapping station network design that will perform well under different realizations of demand.

In this paper, we propose two models to address the problem faced by BP, or other charging service providers following similar strategies. We develop our models by utilizing recent advances in distributionally-robust optimization [20, 21, 29], a framework for making optimal decisions under limited and/or imprecise information. The first model (the "cost-concerned" model) assumes that the service provider is mainly concerned about cost, and minimizes a robust estimate of the expected building and operating costs of the system. The second model (the "goal-driven" model) assumes that the service provider is more concerned about meeting a certain return-on-investment (ROI) target, and attempts to maximize a robust estimate of the probability of doing so.

The remainder of this paper is organized as follows. In Section 2, we first discuss the significance of our research to China. Then, in Section 3, we review the related literature. We then formally introduce the problem setting in Section 4. Next, in Sections 5 and 6, we present the cost-concerned and goal-driven network design models, respectively. Finally, we discuss our findings from computational studies in Section 7 and conclude the paper in Section 8.

2. Significance to China

With a booming economy, China faces the grand challenge of maintaining mobility growth while resulting curbing the environmental footprint. If China can leapfrog gasoline cars with EVs, the environmental benefits will be tremendous. Observing the economic and environmental potential of EVs, the central and local governments in China have taken prompt actions, including a \$15 billion investment in an alliance of 16 large state-owned companies formed to pursue standardization, research and development related to EVs [45]. As a supporting measure, incentives of up to RMB60,000 (\$9,500) per EV purchase are provided in the pilot cities of Shanghai, Changchun,

Shenzhen, Hangzhou and Hefei [17]. It is China’s vision to become the worldwide leader in both EV manufacturing and adoption.

When discussing BP’s innovative business model, many do not notice that the the first country adopting battery swapping for EVs at a practical scale was in fact China. Starting from the 2008 Olympics, the city of Beijing has been using buses with battery swapping technology on a number of routes [42]. Recently, Beijing has announced its plan to deploy 250 battery charging and swapping stations by 2015 [22]. Led by the State Grid in Northern China [23], the partnership between BP and China Southern Power Grid in the South [24], ambitious battery swapping station deployment plans are under way in major cities. Compared to alternative technologies such as DC quick charging, China’s central government has specifically selected battery swapping as the core component of the overall EV development strategy using “battery swap mode as a main, plug-in mode as a secondary” means for charging [17]. Part of the reason for this development is that battery swapping allows the grid operators to maintain control over charging schedules and thus the load on the grid [16]. While battery swapping is widely considered a yet-to-be-proven mode of operations worldwide, China is one of the early adopters that are pushing hard for this technology to succeed. Therefore, the potential impact of our research is likely to be even more significant to China than to most other countries in the world.

3. Literature Review

There has only been a limited number of recent studies on the important problem of locating support infrastructure for alternative fuel vehicles (AFVs). Kang *et al.* [32] study the problem of locating biofuel refineries. Their main focus is to optimize the production and distribution of biofuel, rather than to deploy refueling infrastructure. Bai *et al.* [3] study a related model that incorporates congestion effects within the supply chain.

Kuby and Lim [34] study a flow-refueling location problem for generic AFVs, not specific to any particular type. The main focus is to “capture” as much vehicle flow as possible by covering travel paths so that vehicles can always be refueled before running out of fuel, by locating a given number of refuel stations. In [35], the authors extend the earlier work to allow location of stations along arcs, in addition to nodes, of the network. Upchurch *et al.* [49] extend the model to account for the fact that refueling stations have limited capacity. In a railway context, Nourbakhsh and Ouyang [40] study the joint problem of locating refuel stations in a railroad network and scheduling refuels for trains. They formulate the problem as a mixed-integer linear program and propose an efficient Lagrangian relaxation solution algorithm. These models fall within the general category of flow-intercepting location models [7, 8, 9, 31], in which facilities are located to capture flows in

a network by intercepting them along their respective paths. The refueling problem is particularly difficult because flows may be intercepted (i.e., refueled) multiple times.

The central focus of the above refuel location papers has been on ensuring that vehicles get access to refuel stations before running out of fuel. Although this is clearly one of the important concerns in AFV infrastructure planning, there are many other critical considerations that must also be included in a realistic decision-support model for service providers such as BP. For instance, depending on the type of AFV considered and the nature of the refueling operation, capacity limits should be modeled differently. The capacities of battery swapping stations are constrained by the number of spare batteries stocked at the station and the recharging process. These important characteristics cannot be reflected by a simple upper bound on the vehicle flow rate through the station.

Because service providers have to roll out their infrastructure plans before the majority of customers subscribe to the services, it is critical to make plans based on models that are robust with respect to demand uncertainty. The theory of robust optimization (see, e.g., [5, 6, 10, 11]) may help achieve this objective. In particular, because some limited information, such as the best estimate and the possible range of demand parameters may be already known at the planning stage, it is possible to utilize recent advances in distributionally-robust optimization [20, 21, 29]. One key to using this methodology in practice is to tightly approximate robust optimization problems using computationally-tractable conic formulations. Examples of recent applications include inventory control [47], warehouse management [1] and portfolio optimization [38] under uncertainty.

4. Problem Description and Notation

In our basic problem setting, the charging service provider wishes to locate a number of battery swapping stations at strategic locations along a network of freeways. As discussed before, the location of charging adaptors are relatively easy to determine and are not considered in our current paper. The decision problem consists of two-stages. In Stage 1 (the planning stage), the service provider determines the swapping station locations, with only limited information on hand regarding the distribution (e.g., mean, covariance and support) of the uncertain factors governing demand. Then, at the beginning of Stage 2, the uncertain factors are realized and observed. With the additional information, the firm then stocks a sufficient number of batteries at each station to guarantee certain service requirements. We will now discuss the details of the problem setting below. A summary of all notation used is provided in the Electronic Companion¹.

¹The Electronic Companion can be downloaded with the full version of this paper at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2022651.

4.1. Freeway Network

We consider a freeway network² connecting a number of cities, e.g., that of Israel, the first market BP is entering. There are a set of inter-city travel paths, denote by P , on the network. Each path consists of a number of linked segments, which can possibly be defined as sections between adjacent exits. The service provider is given a set of candidate locations, denoted by J , for swapping stations. For example, BP has signed a deal with Dor Alon, a gas station chain in Israel, to locate swapping facilities at the existing gas stations [12]. In that case, the set J will be the set of existing gas stations of the chain, perhaps plus a number of other strategic locations. We define binary decision variables $X_j, j \in J$ to represent whether a station is built at a candidate location j ($X_j = 1$), or not ($X_j = 0$). Because they represent the location decisions, the X_j variables are Stage 1 decision variables. If a station is located, the service provider incurs an annualized fixed cost of f_j .

All EVs traveling long (round) trips must get access to swapping stations before the batteries run out. Therefore, they must visit at least one station along each portion of travel path that is longer than a certain distance $0.5d_C$, where d_C is the maximum travel range allowed with one full charge. We define “subpaths” as sections of travel paths that, as a collection of connected segments, are just longer than $0.5d_C$. We denote the set of such subpaths longer than $0.5d_C$ by Q . For example, consider a 100-mile path consisting of ten 10-mile segments shown in Figure 1(a), ordered from 1 to 10 in sequence, and let d_C be 100 miles. Then, the first subpath consisting of segments 1 to 5 (50 miles long) must pass through at least one station, and so do the next subpaths consisting of segments 2 to 6, 3 to 7, and so on. This way, no matter where a trip is started, or restarted after a swap, the EV will pass by some station before the battery runs out. It is possible that one subpath is in the overlapping region of several paths. We use the binary parameter b_{pq} to denote whether subpath $q \in Q$ is part of path $p \in P$ ($b_{pq} = 1$) or not ($= 0$). We further define a binary parameter a_{jq} to indicate whether candidate location $j \in J$ is along subpath $q \in Q$ ($a_{jq} = 1$) or not ($a_{jq} = 0$). We further define binary decision variables Z_{jq} to indicate whether EVs traveling along subpath $q \in Q$ will visit a swapping station at location $j \in J$. Likewise, the binary decision variables Y_{jp} indicate whether EVs traveling path $p \in P$ will visit a swapping station at location $j \in J$.

4.2. Basic Demand Model

In order to characterize the operations of swapping stations, we have to first model the rate at which EVs patronize stations for service. As mentioned above, EVs traveling on some subpath

² Because battery swapping is needed for long-distance travel, it is sufficient to consider only freeways and ignore intra-city streets.

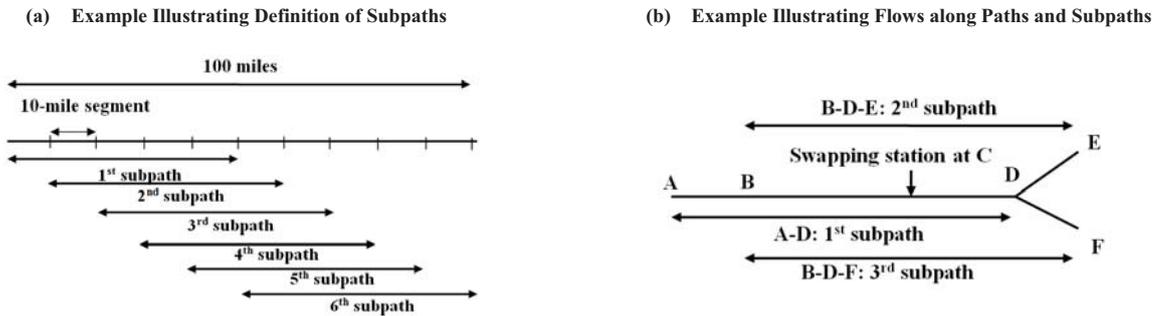


Figure 1 Definition of Paths and Subpaths

$q \in Q$ will visit at least one station along the subpath.³ Because a subpath can lie in the overlapping region of multiple paths, the number of EVs patronizing the station in question will be the sum of flows of all paths covering it. In the example shown in Figure 1(b), consider the two travel paths, from A to E and from A to F. The basic refueling requirements ensure that all EVs traveling the three subpaths, A to D (on which the two paths overlap), B to E and B to F, must visit some station. This is satisfied by locating one station at C. In this case, the demand rate seen at the station at C will be the total flow rate of the two origin-destination paths, A to E and A to F.

We assume that swap-demanding EVs traveling the path $p \in P$ enter the freeway network following a Poisson process with rate λ_p . We consider two types of vehicles requiring battery swaps: (1) all of those traveling along paths longer than the range of the battery d_C , and (2) those traveling round-trips along paths with length between $0.5d_C$ and d_C . Therefore, the rate λ_p will be the total flow of EVs if path p is longer than d_C , and the component of round trip flow if the length of path p is between $0.5d_C$ and d_C . The arrival processes corresponding to different paths are assumed to be mutually independent. Because the swapping demand on different paths can be directed to the same station(s), the demand arrival at a particular swapping station at $j \in J$ will be the superposition of a number of independent Poisson processes (assuming constant travel time to the station). Therefore, with a slight abuse of notation, the arrival process of EVs wishing to swap batteries at the station will be a Poisson process with rate $\lambda_j = \sum_{p \in P} \lambda_p Y_{jp}$.

In Stage 1, as discussed before, the service provider does not know the exact values of the arrival rates. They depend on various factors, such as the adoption rate of EVs, the geographical

³ Here, we assume that all EVs traveling the same path to visit the same stations, but this assumption can be easily relaxed by redefining one path to multiple ones with the same origin and destination, and defining extra subpaths corresponding to flows on these paths for the same freeway section.

characteristics of path p and the driving behavior of EV drivers, many of which are not precisely known in Stage 1. Our model of uncertainty is based on the following assumption:

ASSUMPTION 1. *The path-based demand arrival rate, λ_p , is a linear function of a number of mutually independent random factors, $\tilde{z}_l, l = 1, \dots, L$, known as the “primitive uncertainties”, i.e.,*

$$\lambda_p = \sum_{l=1}^L \hat{\lambda}_{pl} \tilde{z}_l, \quad (1)$$

where $\hat{\lambda}_{pl}, l = 0, \dots, L$ are known constants. In Stage 1, only the means, supports, and variances of the element of the random vector $\tilde{\mathbf{z}} = (\tilde{z}_1, \dots, \tilde{z}_L)$ are known, but not the precise distribution. The family of joint distributions with the given means, variances and supports is denoted by \mathbb{F} , which is assumed to be non-empty. At the beginning of Stage 2, the realization of $\tilde{\mathbf{z}}$ is precisely observed.

The random vector $\tilde{\mathbf{z}}$ may represent factors such as the regional EV adoption rate and the flow of EV traffic along paths. The known means (denoted by $\mu_l, l = 1, \dots, L$), supports (denoted by $[\underline{z}_l, \bar{z}_l], l = 1, \dots, L$) and covariance matrix (denoted by Σ with $\sigma_l, l = 1, \dots, L$ being the individual standard deviations) of the primitive uncertainties represent the best estimate, the estimated range of possible realizations and the variability, respectively. The linear demand model is quite versatile and can be customized by the decision maker by varying the interpretations of the primitive uncertainties and inputting the corresponding support and moment data. In Section 7, we will provide one reasonable example and run computational experiments with the San Francisco Bay Area freeway network.

4.3. Swapping Station Operations

After the stations are built in Stage 1, and realizations of the primitive uncertainties $\tilde{\mathbf{z}}$ are observed at the beginning of Stage 2, the charging service provider will stock batteries at the swapping stations to meet demand. Let I_j denote the number of batteries to be stocked at a station at site $j \in J$. This quantity depends on both the chosen locations and demand realizations, and thus is a Stage 2 recourse decision variable. Because batteries are expensive and have limited usage lives, holding costs are significant. We use h to denote the annualized holding cost per battery. Furthermore, because stocking more batteries at a station necessitates recharging of more batteries in parallel, the service provider should make sure that the number of batteries at a station at j does not exceed the maximum number that can be safely accommodated by the electric grid, denoted by g_j .

Recall that at a swapping station, EVs arrive and request for swapping service following a Poisson process with rate $\lambda_j = \sum_{p \in P} \lambda_p Y_{jp}$, which is a precisely known quantity now in Stage 2. Consider a depleted battery that is unloaded from an incoming EV. The battery will be recharged at the

station until being swapped onto another incoming EV. Assuming that the physical swapping operation is instantaneous⁴, there are always I_j batteries stored at the station at any moment of time. For tractability reasons, we further assume that the unloaded batteries are reused in a first-in-first-out (FIFO) order. Note that this policy is not necessarily optimal, because incoming batteries arrive with varying states-of-charge (SOC), such that the one that arrives the earliest (and having been recharged for the longest) is not necessarily the fullest one or the most suitable one to use for the incoming EV. Nevertheless, we conjecture that FIFO is a good heuristic policy due to the physical nature of battery recharging. Typical constant current-constant voltage (CC-CV) charging technologies lead to charging curves with initial periods of sharp increases in SOCs followed by longer durations of reduced-rate charging (e.g., [18]). As battery swapping is intended to be a range extension measure and a secondary option for refueling (versus the primary mode of home overnight recharging), there is a high chance that the interarrival times between consecutive EVs (and depleted batteries) are at least several minutes. Then, during these few minutes where the battery arriving earlier enjoys the highest rate of recharge, much of the variability of the initial SOCs will be eliminated.

To validate this conjecture, we perform the following simulation experiment using realistic parameter values. We use the most intuitive implementable policy, highest-SOC-first (HSF), as a benchmark for comparison. Under the HSF policy, an arriving EV will be given the battery in inventory that has the highest SOC, i.e., the closest to being fully-recharged. In the experiment, we consider a swapping station facing Poisson EV arrivals with rates of 15 EVs per hour. SOCs are modeled on a scale from 0 (empty) to 1 (full), where incoming batteries have random SOCs (denoted by β) drawn from a Normal distribution with mean 0.3 and standard deviation 0.1, truncated in the interval $[0, 1]$. To capture the initial short periods of sharp increases in SOCs followed by longer durations of reduced-rate charging (e.g., [18]), we model the charging curve with an exponential function $\theta(t, \beta) = 1 - (1 - \beta)e^{-2t}$, where $\theta(t, \beta)$ denotes the SOC of a battery that has been recharged for t hours from an initial SOC of β . For example, a completely depleted battery ($\beta = 0$) can be recharged to $\theta = 0.63$ in 30 minutes, $\theta = 0.86$ in 1 hour and $\theta = 0.95$ in 1.5 hours.

In Figure 2, we plot the cumulative distribution functions of the SOCs of batteries picked up by customers, following the FIFO and HSF policies, when the swapping station stocks 5, 15 and 25 batteries. While the battery SOCs following HSF is always higher (i.e., customers pick up closer-to-full batteries), we observe that the differences are minimal. In the realistic cases where there are sufficient numbers (15 to 25) of batteries at the station, the differences between the two policies are negligible. This validates the use of FIFO as a reasonable approximation of the battery stocking requirements, or even as a heuristic operating policy.

⁴The swapping can be completed in 1-2 minutes [46], which is negligible compared to the amount of time needed to fully recharge the battery.

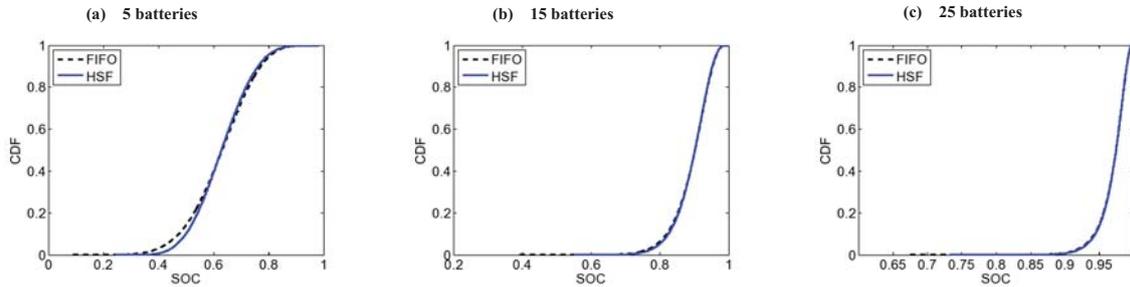


Figure 2 Distributions of Battery SOC's Picked up by Customers Following FIFO and HSF Policies

Under the FIFO policy, each battery, after being unloaded from an incoming EV, stays at the station for recharge until all the other $I_j - 1$ batteries that came before it are used, and the next EV arrives. Because EVs arrive and pick up batteries following a Poisson process, this duration is equal to I_j interarrival times, or the sum of I_j exponential random variables.

BP has announced that service subscribers will be guaranteed access to an inventory of batteries with a guaranteed service level agreement [13]. In this paper, we assume that the service provider sets a service requirement in the following form.

ASSUMPTION 2. It is required that in at least $\alpha (> 0.5)$ proportion of the swapping requests, the EV picks up a battery that has been recharged for at least t time units.

For example, the firm may require that at least 90% of EVs users will pick up batteries that have been recharged for 2 hours or more at the station, or at least 95% pick up batteries that have been recharged for 1 hour or more. Requirements in this form are equivalent to requiring that the number of Poisson demand arrivals in a duration of t time units being less than I_j with probability α . To obtain tractable constraints, we approximate the Poisson demand with a Normal distribution with matching mean and variance. This approximation is commonly used in inventory control (e.g., Shang and Song [48] used it to approximate Poisson demand in a multi-echelon inventory system even with low demand rates), and is particularly accurate when the demand rate is moderate to large (e.g., requiring 10 or more batteries). To satisfy the service requirement, the number of batteries needed is roughly $t\lambda_j + \Phi^{-1}(\alpha)\sqrt{t\lambda_j}$, where $\Phi^{-1}(\cdot)$ is the inverse standard normal CDF. Therefore, the minimum number of batteries needed is given by:

$$I_j = t\lambda_j + \Phi^{-1}(\alpha)\sqrt{t\lambda_j} = t \sum_{p \in P} \lambda_p Y_{jp} + \Phi^{-1}(\alpha) \sqrt{t \sum_{p \in P} \lambda_p Y_{jp}}. \quad (2)$$

With the above problem setting and assumptions, we are now ready to introduce the two optimization models in the next section.

5. Cost-Concerned Model

In this section, we discuss the model that minimizes the expected two-stage cost. This includes the fixed costs of opening and operating the swapping stations plus the expected (with respect to demand realizations) battery holding costs. The basic model that the service provider has to solve in Stage 1 can be formulated as:

$$\min \sum_{j \in J} \left\{ f_j X_j + h \sup_{\mathbb{P} \in \mathbb{F}} E_{\mathbb{P}} [I_j(\mathbf{Y})] \right\} \quad (3)$$

subject to:

$$Y_{jp} \geq b_{pq} Z_{jq}, \text{ for each } j \in J, p \in P, q \in Q \quad (4)$$

$$\sum_{j \in J} a_{jq} Z_{jq} \geq 1, \text{ for each } q \in Q \quad (5)$$

$$Y_{jp} \leq X_j, \text{ for each } j \in J, p \in P \quad (6)$$

$$\inf_{\mathbb{P} \in \mathbb{F}} P_{\mathbb{P}}(I_j(\mathbf{Y}) \leq g_j) \geq 1 - \epsilon_g, \text{ for each } j \in J \quad (7)$$

$$X_j \in \{0, 1\}, \text{ for each } j \in J$$

$$Y_{jp} \in \{0, 1\}, \text{ for each } j \in J, p \in P \quad (8)$$

$$Z_{jq} \in \{0, 1\}, \text{ for each } j \in J, q \in Q. \quad (9)$$

In the above formulation, the objective (3) is to minimize the cost of opening swapping stations and equipping them with swapping machinery and batteries. The number of batteries to stock at a location is a function of the demand routing pattern, i.e., the \mathbf{Y} matrix, as well as the realization of primitive uncertainties $\tilde{\mathbf{z}}$. We assume that the firm is risk neutral and considers the expected cost. Moreover, because the exact distribution of $\tilde{\mathbf{z}}$, denoted by \mathbb{P} , is not precisely known, we apply the notion of distributionally-robust optimization [20, 21, 29] and consider the expected cost under the worst possible (i.e., supremum) distribution over the family \mathbb{F} of all possible distributions satisfying the mean (μ_l), variance (σ_l^2) and support ($[\underline{z}_l, \bar{z}_l]$) information, as well as independence of \tilde{z}_l 's. Recall that we assume the family \mathbb{F} to be non-empty, i.e., for each l , the given mean, variance and support can be achieved by some univariate distribution.

Constraints (4) suggest that routing flows along a subpath to a particular station is equivalent to directing the flows of all travel paths containing this subpath to the station. Constraints (5) make sure that EVs traveling a subpath longer than the tolerance distance will visit at least one station along the subpath. Constraints (6) require that a station must be opened at locations to which swapping demand is directed. Finally, the chance constraints (7) require that with a (high) probability of at least $1 - \epsilon_g$, the number of batteries at a station does not exceed the maximum allowable number that can be recharged in parallel by the local electric grid. Note that

the constraints act on the worst-case (infimum) probability over all possible distributions in family \mathbb{F} .

In the remainder of this section, we will focus on how to find a tight upper bound (i.e., slightly conservative) approximation to the problem by solving a mixed-integer second-order cone program (MISOCP). Commercial integer programming solver packages such as CPLEX include advanced branch-and-bound algorithms capable of solving large-scale MISOCPs. Therefore, by approximating the original problem with the upper bound problem in MISOCP form, we will be able to solve practical-sized instances of it using commercial solvers, as we demonstrate in Section 7.

There are two key steps in this approximation process: to find a tight upper bound on the worst-case expected value of the battery cost in (3) and to come up with a tractable but tight restriction on the chance constraints on battery numbers at swapping stations (7). We discuss these two steps in Sections 5.1 and 5.2, respectively.

5.1. Distributionally-Robust Bound on Expected Cost

The number of batteries to stock at a station is a function of the realized rate of demand routed to the station, as given by equation (2). To facilitate the analysis, we first transform the problem slightly by introducing new decision variables V_j to represent the worst-case expected number of batteries at location $j \in J$.

LEMMA 1. *The cost-concerned problem (3-9) is equivalent to:*

$$\min \sum_{j \in J} (f_j X_j + h V_j) \quad (10)$$

subject to:

$$V_j \geq \sup_{\mathbb{P} \in \mathbb{F}} E_{\mathbb{P}} \left[t \sum_{p \in P} \lambda_p Y_{jp} + \Phi^{-1}(\alpha) \sqrt{t \sum_{p \in P} \lambda_p Y_{jp}} \right], \text{ for each } j \in J \quad (11)$$

plus (4-9).

Then, we present the following result that provides the worst-case expected value on the right hand side of (11). The proof of all analytical results are provided in the Electronic Companion.

PROPOSITION 1. *For $l = 1, \dots, L$, define independent univariate random variables \tilde{z}_l following two-point distributions with support points $\left(\frac{\mu_l \tilde{z}_l - (\mu_l^2 + \sigma_l^2)}{\tilde{z}_l - \mu_l}, \tilde{z}_l \right)$, with respective probabilities $\left(\frac{(\tilde{z}_l - \mu_l)^2}{(\tilde{z}_l - \mu_l)^2 + \sigma_l^2}, \frac{\sigma_l^2}{(\tilde{z}_l - \mu_l)^2 + \sigma_l^2} \right)$.*

Then, the supremum expected value in (11) is given by:

$$\sup_{\mathbb{P} \in \mathbb{F}} E_{\mathbb{P}} \left[t \sum_{p \in P} \lambda_p Y_{jp} + \Phi^{-1}(\alpha) \sqrt{t \sum_{p \in P} \lambda_p Y_{jp}} \right] = t \sum_{p \in P} \sum_{l=1}^L \hat{\lambda}_{pl} \mu_l Y_{jp} + \Phi^{-1}(\alpha) \sqrt{t E \left[\sum_{p \in P} \sum_{l=1}^L \hat{\lambda}_{pl} Y_{jp} \tilde{z}_l \right]}. \quad (12)$$

Proposition 1 gives the exact value of the right hand side of constraints (11). However, it is not easy to use this result directly in an optimization problem, because it is not in a convenient form that can be handled by most commercial solvers. The next result allows us to obtain tight upper and lower bounds on this expression in integer SOCP form, which can be handled by commercial solvers such as CPLEX.

PROPOSITION 2. *Suppose $\hat{\lambda}_{pl}$ is nonnegative for each $p \in P, l = 1, \dots, L$. Moreover, \tilde{z}_l has non-negative support for each $l = 1, \dots, L$. Let $a_l = \frac{\tilde{z}_l}{\mu_l}$ and $b_l = \frac{\sigma_l}{\mu_l}$ for $l = 1, \dots, L$. Let $a = \min_{l=1, \dots, L} \{a_l\}$, $b = \max_{l=1, \dots, L} \{b_l\}$, $a' = \max_{l=1, \dots, L} \{a_l\}$ and $b' = \min_{l=1, \dots, L} \{b_l\}$ for $l = 1, \dots, L$. Then, the following bound upper bound holds:*

$$\sup_{\mathbb{P} \in \mathbb{F}} E_{\mathbb{P}} \left[t \sum_{p \in P} \lambda_p Y_{jp} + \Phi^{-1}(\alpha) \sqrt{t \sum_{p \in P} \lambda_p Y_{jp}} \right] \leq t \sum_{p \in P} \sum_{l=1}^L \hat{\lambda}_{pl} \mu_l Y_{jp} + \Phi^{-1}(\alpha) \sqrt{t} \bar{\Psi} \sqrt{\sum_{p \in P} \sum_{l=1}^L \hat{\lambda}_{pl} \mu_l Y_{jp}^2};$$

furthermore, if $a \geq b^2 + 1$, the following lower bound holds:

$$\sup_{\mathbb{P} \in \mathbb{F}} E_{\mathbb{P}} \left[t \sum_{p \in P} \lambda_p Y_{jp} + \Phi^{-1}(\alpha) \sqrt{t \sum_{p \in P} \lambda_p Y_{jp}} \right] \geq t \sum_{p \in P} \sum_{l=1}^L \hat{\lambda}_{pl} \mu_l Y_{jp} + \Phi^{-1}(\alpha) \sqrt{t} \underline{\Psi} \sqrt{\sum_{p \in P} \sum_{l=1}^L \hat{\lambda}_{pl} \mu_l Y_{jp}^2},$$

where:

$$\underline{\Psi} = \sqrt{a} - \frac{a-1}{\sqrt{a} + \sqrt{1 - \frac{b^2}{a-1}}}, \quad \bar{\Psi} = \sqrt{a'} - \frac{a'-1}{\sqrt{a'} + \sqrt{1 - \frac{b'^2}{L(a'-1)}}}.$$

Note that the maximum relative error of using the upper bound, by comparing with the lower bound, is $\bar{\Psi}/\underline{\Psi} - 1$, which depends only on input parameters, but not the decision variables. We evaluate the tightness of the upper and lower bounds with a numerical experiment. In this experiment, we consider $L = 10$ and randomly generate μ_l , a_l and b_l from uniform distributions on intervals $[0, 1]$, $[2, 3]$ and $[0, 1]$, respectively. For each sample of distributional parameters generated, we calculate the exact supremum of the expected value by Proposition 1 and the bounds in Proposition 2. In all 100 instances generated, relative errors of the upper bound values are less than 1%, and those of the lower bound range from about 5% to 20%. This suggests that the upper bound, which will be used in the subsequent model development and computational experiments, is very accurate.

Finally, we note that the Y_{jp} variables inside the square root term in the upper bound function are squared. Because the Y_{jp} variables are binary, squaring them does not change their values. We do so because it allows us to express the upper bound in integer SOCP form. Replacing the right hand side of constraints (11) by the upper bound, we obtain the new set of (slightly) restricted constraints:

$$V_j \geq \sum_{p \in P} \sum_{l=1}^L t \hat{\lambda}_{pl} \mu_l Y_{jp} + \Phi^{-1}(\alpha) \bar{\Psi} \sqrt{\sum_{p \in P} \sum_{l=1}^L t \hat{\lambda}_{pl} \mu_l Y_{jp}^2}, \text{ for each } j \in J. \quad (13)$$

The constraints (13) are in SOCP form. Moreover, any solution that satisfies them will also satisfy the original set of constraints (11). This completes the approximation of the worst-case expected battery costs using SOCP form. The next step is to do the same for the chance constraints (7).

5.2. Chance Constraints on Batteries

Recall that the number of batteries required is given by (2), which is a strictly increasing function in the demand rate λ_j . This suggests that constraining on the number of batteries is equivalent to constraining on the demand rate, as suggested by the following result.

PROPOSITION 3. *The chance constraints (7) are equivalent to:*

$$\inf_{\mathbb{P} \in \mathbb{F}} P_{\mathbb{P}} \left(\sum_{p \in P} \sum_{l=1}^L \hat{\lambda}_{pl} \tilde{z}_l Y_{jp} \leq \hat{g}_j \right) \geq 1 - \epsilon_g \text{ for each } j \in J, \\ \text{where } \hat{g}_j = \frac{\left(\sqrt{g_j + \Phi^{-1}(\alpha)^2/4} - \Phi^{-1}(\alpha)/2 \right)^2}{t} \text{ is a constant.} \quad (14)$$

With this result, we have transformed the original chance constraints on nonlinear expressions for the number of batteries into ones that act on demand rates, which are linear functions of both the primitive uncertainties and the decision variables. It is known that tight restrictions on such chance constraints are available in the form of more tractable conditional value-at-risk (CVaR) constraints [39]. Tight bounds on the distributionally-robust CVaR constraints in SOCP form are also derived in [21, 19]. Applying these results, we are then able to tightly approximate our original problem as an MISOCP. Furthermore, solving the MISOCP gives a feasible solution to the original problem and an upper bound on its optimal objective value. The complete formulation of the MISOCP can be found in the Electronic Companion.

5.3. Solution Accuracy

The upper bound developed in Proposition 2 allows us to obtain a tractable MISOCP formulation to be used in the computational studies. Before proceeding, we validate the accuracy of this tractable formulation by comparing its output against sample-average approximations (SAA). To perform such comparisons, we consider the cases where the \tilde{z}_l random variables follow Normal (with the same means and variances as we input to the robust models), uniform (with the same means and variances) and triangular distributions (with the same means and supports). We note

that SAA-based formulations for the swapping station location model, with reasonable sample size (10,000 in our tests), are computationally intractable. Therefore, to draw comparisons, we perform comparisons *for randomly generated values of the \mathbf{X}, \mathbf{Y} variables*, i.e., we randomly generate solutions and evaluate them using two alternative methods. If the two sets of objective evaluations are perfectly associated, then our robust objective function provide accurate approximations to the SAA formulations. The \mathbf{X} values are randomly set to 1 with a probability equal to the proportion of X_j variables equal to 1 in the optimal solution to the robust formulation. Similarly, within the sites $j \in J$ where X_j is set to 1, we randomly set $Y_{jp} = 1$ with a probability equal to the proportion of Y_{jp} variables equal to 1 out of those $j \in J$ with $X_j = 1$ in the optimal solution to the robust formulation. Because the fixed location costs and the linear component of the battery cost expressions are the common under both the robust model and SAA, they are excluded in the objective evaluations for comparison.

Our results show that our formulation performs extremely well. Out of 5,000 random solutions generated, the average relative error for all Normal, uniform and triangular distributions are 2.1%, 1.1% and 2.8%, respectively. More importantly, the (Spearman) rank correlation between the robust evaluations and the worst objective across the three distributions is over 0.99. Therefore, the robust model ranks the set of solutions in an ordering that almost completely coincides with the ranking based on the worst objective value of the three distributions. This shows the remarkable accuracy of our model. In the next section, we proceed to discussing the alternative goal-driven model.

6. Goal-Driven Model

BP is a start-up backed by investments from investors such as HSBC and Morgan Stanley [45]. The starting capital, considering its ambitious plans, is constrained. It is also natural for its investors to set some target on profit or return-on-investment (ROI) rate, possibly by benchmarking against predicted performances of alternative investment opportunities, such as stock market returns. Therefore, besides minimizing expected cost, an alternative objective that makes sense for BP, and other similar service providers, is to maximize the likelihood of meeting such a target.

Although battery swapping is one of the crucial elements in the business model, its role is to provide refueling coverage to users, instead of being a key revenue driver. In the service contract, users are guaranteed convenient access to swapping service as part of a complete service package. Therefore, we assume that the service provider's revenue is exogenous to the swapping station location decisions. Under this assumption, the location decisions only affect the firm's profit through cost.

Considering the limitation on available funds, we assume that the firm has a fixed budget B that it can spend in Stage 1 to construct the swapping stations. We assume that the variable profit,

i.e., portion of profit that changes with the realized demand (unknown in Stage 1), is given by an exogenous variable profit, which includes the revenue from subscription and usage minus all costs incurred outside of the swapping infrastructure system, minus the battery holding costs at swapping stations. In the goal-driven model, we assume that the firm's goal is to make sure that a given target ROI rate T is reached. The ROI is defined to be the annual variable profit (consisting of the exogenous portion minus the battery holding costs) divided by the Stage 1 fixed investments on building the swapping stations (i.e., $\sum_{j \in J} f_j X_j$). This assumption can be made even if the overall profit contains a fixed and deterministic portion, because then we may simply subtract the fixed portion from the target in advance.

For simplicity, we assume that the exogenous portion of the variable profit is given by an affine function of realized path flow rates (λ_p). This is a reasonable assumption considering that the firm's revenue should be proportional to the total number of vehicle-mile travelled. Cost excluding those related to swapping mainly come from the cost of electricity, which is proportional to mileage, and the cost of building charging slots, which is roughly proportional to the number of users because charging slots will be deployed at the users' homes and work places. With the above, we may express the exogenous part of the variable profit as $\pi_0 \lambda_0 + \sum_{p \in P} \pi_p \lambda_p$, where $\pi_p, p \in P \cup \{0\}$ are the marginal exogenous (annual) variable profits per unit flow on the associated paths, which can be calculated based on both the lengths of the travel paths and a profit margin per vehicle-mile travelled.

Then, the goal-driven swapping station location model can be formulated as follows:

$$\max_{\mathbb{P} \in \mathbb{F}} \inf_{\mathbb{P} \in \mathbb{F}} P_{\mathbb{P}} \left(\pi_0 \lambda_0 + \sum_{p \in P} \pi_p \lambda_p - \sum_{j \in J} h I_j \geq T \sum_{j \in J} f_j X_j \right) \quad (15)$$

subject to:

$$\sum_{j \in J} f_j X_j \leq B \quad (16)$$

$$Y_{jp} \geq b_{pq} Z_{jq}, \text{ for each } j \in J, p \in P, q \in Q \quad (17)$$

$$\sum_{j \in J} a_{jq} Z_{jq} \geq 1, \text{ for each } q \in Q$$

$$Y_{jp} \leq X_j, \text{ for each } j \in J, p \in P$$

$$\inf_{\mathbb{P} \in \mathbb{F}} P_{\mathbb{P}} (I_j(\mathbf{Y}) \leq g_j) \geq 1 - \epsilon_g, \text{ for each } j \in J \quad (18)$$

$$X_j \in \{0, 1\}, \text{ for each } j \in J$$

$$Y_{jp} \in \{0, 1\}, \text{ for each } j \in J, p \in P \quad (19)$$

$$Z_{jq} \in \{0, 1\}, \text{ for each } j \in J, q \in Q. \quad (20)$$

The objective (15) is to maximize the distributionally-robust estimate of the probability that the ROI target T is met. The constraint (16) makes sure that the Stage 1 cost incurred to build the stations does not exceed the available budget. The remaining constraints (17-20) are directly inherited from the cost-concerned model discussed in Section 5. We have already discussed how to approximate the chance constraints (18) in SOCP form in Section 5.2. We now focus on how to transform the problem so that the chance-based objective function can be handled by solving a sequence of MISOCP instances.

6.1. Approximation and Transformations

The approximations and transformations to be outlined in this section are based on the results of [19]. We first use the following simple result to transform the problem into one that involves a chance constraint, with the probability level being a decision variable as well.

LEMMA 2. *The goal-driven problem (15) is equivalent to:*

$$\max 1 - \gamma \tag{21}$$

subject to:

$$\inf_{\mathbb{P} \in \mathbb{F}} P_{\mathbb{P}} \left(\pi_0 \lambda_0 + \sum_{p \in P} \pi_p \lambda_p - \sum_{j \in J} h I_j \geq T \sum_{j \in J} f_j X_j \right) \geq 1 - \gamma \tag{22}$$

plus (16-20).

If γ is fixed, the goal-driven model is equivalent to one with a chance constraint on the total number of batteries at all the swapping stations. However, such a chance constraint is not analytically tractable, especially that I_j are nonlinear functions in both the primitive uncertainties and the decision variables, given by (2). As suggested by [19], we approximate chance constraints by using a more tractable risk measure known as the conditional value-at-risk (CVaR). As proven by [39], the CVaR measure provides the tightest possible bound that preserves convexity for a linear chance constraint. Even though our chance constraint (22) is nonlinear, we will see that using a CVaR constraint still allows us to obtain a much more tractable approximation to the original problem. Note that after such an approximation, the resulting objective will be to maximize the “shortfall aspiration level criterion” proposed by [19]. The following result on the relationship of the new and original problems will also hold.

LEMMA 3. *If the constraint (22) is replaced by the following constraint on the conditional value-at-risk:*

$$\sup_{\mathbb{P} \in \mathbb{F}} \psi_{1-\gamma} \left(-\pi_0 \lambda_0 - \sum_{p \in P} \pi_p \lambda_p + \sum_{j \in J} h I_j + T \sum_{j \in J} f_j X_j \right) \leq 0, \tag{23}$$

$$\text{where } \psi_{1-\gamma}(\tilde{v}) = \min_{\theta} \left(\theta + \frac{1}{\gamma} E_{\mathbb{P}}(\tilde{v} - \theta)^+ \right),$$

then the optimal solution to the new problem gives a feasible solution to the original problem, and the optimal objective value (i.e., the value of γ), provides a lower bound on the corresponding objective value given by this solution in the original problem.

Our next step is to obtain tight bounds on the constraints (23). Note that by substituting the required number of batteries as given by (2), we obtain the following constraint:

$$\sup_{\mathbb{P} \in \mathbb{F}} \psi_{1-\gamma} \left(-\pi_0 \sum_{l=1}^L \hat{\lambda}_{0l} \tilde{z}_l - \sum_{p \in P} (\pi_p - ht) \sum_{l=1}^L \hat{\lambda}_{pl} \tilde{z}_l + h\Phi^{-1}(\alpha) \sum_{j \in J} \sqrt{t \sum_{p \in P} \sum_{l=1}^L \hat{\lambda}_{pl} \tilde{z}_l Y_{jp} + T \sum_{j \in J} f_j X_j} \right) \leq 0. \quad (24)$$

For goal-driven problems with objective functions involving expressions linear in the decision variables and the primitive uncertainties, Chen and Sim [19] propose a solution approach based on searching for the value of γ such that the problem with CVaR objective has an optimal objective of 0, i.e., identifying the γ level at which the CVaR constraint is binding. However, the constraint (24) involves a nonlinear (square root) expression. To provide a tractable approximation to this constraint, we use the following result, which is obtained by using a collection of linear functions to bound the square root function from above.

PROPOSITION 4. *The following CVaR expression, involving an expression linear in the primitive uncertainties, is an upper bound of the CVaR value in (24):*

$$\sup_{\mathbb{P} \in \mathbb{F}} \psi_{1-\gamma} \left(-\pi_0 \sum_{l=1}^L \hat{\lambda}_{0l} \tilde{z}_l - \sum_{p \in P} (\pi_p - ht) \sum_{l=1}^L \hat{\lambda}_{pl} \tilde{z}_l + h\Phi^{-1}(\alpha) t \sum_{j \in J} \frac{\eta_j}{2} \sum_{p \in P} \sum_{l=1}^L \hat{\lambda}_{pl} \tilde{z}_l Y_{jp} \right) + \sum_{j \in J} \frac{h}{2\eta_j} + T \sum_{j \in J} f_j X_j, \quad (25)$$

where $\eta_j, j \in J$ can be any positive numbers. In particular, given fixed binary values of \mathbf{X} and \mathbf{Y} , one can obtain the tightest possible bound in such form by using the optimal η values that solves the following problem:

$$\min_{\eta_j \geq 0, j \in J} \sup_{\mathbb{P} \in \mathbb{F}} \psi_{1-\gamma} \left(-\pi_0 \sum_{l=1}^L \hat{\lambda}_{0l} \tilde{z}_l - \sum_{p \in P} (\pi_p - ht) \lambda_p + \sum_{j \in J} h\Phi^{-1}(\alpha) t \frac{\eta_j}{2} \sum_{p \in P} \lambda_p Y_{jp} \right) + \sum_{j \in J} \frac{h}{2\eta_j} + T \sum_{j \in J} f_j X_j. \quad (26)$$

By selecting different η values, we may generate a collection of tight upper bounds by applying Proposition 4. These CVaR expressions can be bounded in SOCP form using the results in Chen and Sim [19]. Suppose we identify N of these bounds by specifying η_{jn} values for $n = 1, \dots, N$. Then, we may apply the following result to obtain a unified bound:

PROPOSITION 5. *An upper bound on the worst-case CVaR in (24) can be obtained by solving the following optimization problem:*

$$\min \sum_{n=1}^N \hat{\psi}_n + T \sum_{j \in J} f_j X_j \quad (27)$$

subject to (16-20), and:

$$\begin{aligned} \hat{\psi}_n &= \sup_{\mathbb{P} \in \mathbb{F}} \psi_{1-\gamma} \left(\sum_{l=1}^L \hat{c}_{ln} \tilde{z}_l + h \Phi^{-1}(\alpha) t \sum_{j \in J} \frac{\eta_{jn}}{2} \sum_{l=1}^L \tilde{z}_l \hat{y}_{jln} \right) + \sum_{j \in J} \frac{1}{2\eta_{jn}} \hat{d}_{jn} \\ \sum_{n=1}^N \hat{c}_{ln} &= -\pi_0 \hat{\lambda}_{0l} - \sum_{p \in P} (\pi_p - ht) \hat{\lambda}_{pl}, \text{ for each } l = 1, \dots, L \\ \sum_{n=1}^N \hat{d}_{jn} &= h, \text{ for each } j \in J \\ \sum_{n=1}^N \hat{y}_{jln} &= \sum_{p \in P} \hat{\lambda}_{pl} Y_{jp}, \text{ for each } j \in J, l = 1, \dots, L \\ \hat{y}_{jln} &\geq 0, \text{ for each } j \in J, l = 1, \dots, L, n = 1, \dots, N. \end{aligned}$$

Furthermore, this unified bound is tighter than the minimum of the N individual upper bounds.

Proposition 5 allows us to replace the complex CVaR constraint (24) by one that involves the CVaR values of a collection of linear expressions. We may then use the tight bounds proposed by [19] to approximate the problem. Combining Propositions 4 and 5, we may perform an iterative procedure to generate tight bounds by solving problem (26) in Proposition 4 (approximated in SOCP form) given some $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ values, improve the unified bound in Proposition 5, and obtain new $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ values. This process can be repeated until a certain stopping criterion is met. This procedure is similar to the piecewise-linear approximation procedure for a deterministic objective value with square root terms proposed by Magnanti *et al.* [36].

We have discussed how to approximate the constraint (22) by a more tractable one, if γ is known and fixed. Unfortunately, as also noted by [19], the problem is not jointly convex in γ and the other decision variables. To optimize over γ and solve problem (21), we may adapt the binary search procedure proposed by [19] by solving problem (27) iteratively.

ALGORITHM 1.

1. Set $\underline{\gamma} := 0$ and $\bar{\gamma} := 1$.
2. Set $N = 1$, $\eta_{j1} := 1$ for all $j \in J$. If $\bar{\gamma} - \underline{\gamma} < \zeta$ (some pre-specified precision level), stop and output the incumbent solution.
3. Let $\gamma := 0.5(\bar{\gamma} + \underline{\gamma})$.
4. Solve problem (27) by using $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ as decision variables. Record the optimal solution as the incumbent.

	$\gamma = 0.05$	$\gamma = 0.25$	$\gamma = 0.5$	$\gamma = 0.75$	$\gamma = 0.95$
Normal	1.3%	0.8%	4.6%	0.8%	2.8%
Uniform	1.5%	0.8%	4.5%	0.7%	2.7%
Triangular	1.4%	1.0%	5.1%	1.5%	1.1%

Table 1 Percentage Errors of CVaR Evaluations with Goal-Driven Formulation against SAA

5. Increment N by 1. Fixing $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ to the incumbent values, solve problem (26) by using $\{\eta_j, j \in J\}$ as decision variables. This involves solving the continuous SOCP in the *Electronic Companion*. Denote the optimal η_j values by $\{\eta_j^*\}$ and the optimal objective value by $\hat{\gamma}$. If $|\hat{\gamma} - \gamma| < \xi$ (some pre-specified precision level) and $N \leq \bar{N}$ (inner iteration limit), go to Step 6. Otherwise, let $\eta_{jN} := \eta_j^*$ and go to Step 4.

6. If $Z(\gamma) \leq 0$, update $\bar{\gamma} = \gamma$. Otherwise, update $\underline{\gamma} = \gamma$. Go to Step 2.

Note that Algorithm 1 consists of an outer loop, in which the value of γ is updated, and an inner loop, in which the piecewise-linear approximation of the square root function is improved. In each iteration of the inner loop, it is clear that the CVaR objective value decreases (i.e., the solution becomes more favorable) as the piecewise linear approximation improves. Given sufficient number of iterations, the inner loop will converge to some local optimal solution, from which the linear approximation cannot be further improved. Furthermore, the outer loop is a bisection search over γ and is guaranteed to converge, given any specified level of precision (i.e., ζ).

6.2. Solution Accuracy

As done for the cost-concerned model, we went through a number of transformations and approximations to obtain a tractable MISOCP formulation based on Algorithm 1. We validate the accuracy of our approximations using a similar experiment as done in Section 5.3. As our approximations for the goal-driven model concerns the approximation of the CVaR value (using the linear upper bound and the results from [19]), we compare the CVaR evaluations (for each $\gamma = 0.05, 0.25, 0.5, 0.75, 0.95$) associated with randomly generated solutions using our approximate formulation and SAA, with the common fixed cost component excluded. Again, the comparison shows remarkable accuracy of our approximations. Out of 5,000 random solutions generated, the average relative errors are very small (Table 1), even by setting the inner iteration limit $\bar{N} = 1$. The (Spearman) rank correlations between the robust evaluations and the worst objective for all cases are above 0.99. After confirming the accuracy of our approximations, we proceed to discuss further computational studies.

7. Computational Studies

In this section, we discuss the results of our computational studies performed based on the cost-concerned and goal-driven models and the managerial insights that can be drawn from them. For

demonstration purposes, we use the major freeway network of the San Francisco Bay Area in our computational studies [25]. The details of the demand model used and the corresponding data input are summarized in the Electronic Companion.

7.1. Comparisons of the Two Models

In the first set of computational experiments, we draw comparisons between the goal-driven and cost-concerned models. We solve both the goal-driven and cost-concerned models for every instance generated. We further compute the worst-case expected cost for the goal-driven solution, and the goal-driven objective value for the cost-concerned solution.

The results are summarized in Table 2. The “CC” columns show the results from solving the cost-concerned model, the “GD1” and “GD2” columns show the results from solving the goal-driven model, with two different levels of marginal exogenous variable profit π_p values. The π_p values in the “GD2” instances are equal to two times the corresponding values in the “GD1” instances. The purpose for contrasting these two is to compare the profitability of battery swapping in regions with different business environments. For example, BP’s recently-announced pricing levels for the Denmark market [13] correspond to a usage cost roughly equal to the current gasoline price in Denmark. The reason that BP is able to charge such a high margin over the cost of electricity is that the after-tax purchase cost the executive class EV included in the package is roughly equal to that of a compact car. This purchase cost differential arises from the high vehicle registration tax, which is waived for EVs. In markets like Denmark, BP may obtain higher margins, and higher π_p values. For other markets such as the US, such significant tax benefits may not be available, and BP will have to induce adoptions by offering significant savings in terms of usage costs, resulting in lower margins and lower values of π_p .

We generate problem instances by varying the expected flow volume of EVs *requiring swaps* in the network, from 50 EVs per hour to 1400 EVs per hour. Note that although these numbers may appear small relative to the total flow of traffic (typically in the order of thousands per freeway link), they in fact represent scenarios with significant levels of EV adoption. This is because the majority of (about 90%) of trips made by typical drivers are well short of the maximum travel range in one charge [43], meaning that the total flow of EVs in the whole network will be an order of magnitude higher than the flow of EVs *requiring battery swaps*. For each instance, we report the optimal objective values based on the cost-concerned and goal-driven objectives and the number of swapping stations built. All instances are solved using CPLEX 12.1 on a Dell Precision T7500 Workstation an Intel Xeon X5680 CPU and 48GB of memory, running on Windows 7. For every call of CPLEX, we set the termination criterion to be (i) optimality gap reaching 1% or below, (ii) CPU time reaching 4 days, or (iii) the branch-and-bound algorithm reaches 30,000 nodes.

Exp. Flow	CC					GD1			GD2		
	Stations	Cost	γ_1^*	γ_2^*	Time	γ^*	Cost	Time	γ^*	Cost	Time
50	7	427,504	0.00	0.054	5.5 hr	0.00	428,058	214 s	0.055	429,341	198 s
100	7	485,963	0.033	0.173	195 s	0.035	486,341	107 s	0.176	486,915	130 s
200	7	590,884	0.138	0.251	3.7 hr	0.138	592,693	226 s	0.251	592,693	2166 s
400	8	846,514	0.200	0.423	2.3 hr	0.200	846,308	1102 s	0.426	847,418	478 s
600	9	1.1001×10^6	0.223	0.889	3.6 hr	0.223	1.1001×10^6	1.3 hr	0.889	1.1007×10^6	1.2 hr
800	11	1.4604×10^6	0.230	0.995	96 hr	0.230	1.4058×10^6	886 s	0.996	1.4036×10^6	1486 s
1000	13	1.7352×10^6	0.233	1.00	61 hr	0.230	1.7154×10^6	901 s	1.00	1.7154×10^6	1885 s
1200	16	2.1344×10^6	0.231	1.00	59 hr	0.231	2.0805×10^6	1250 s	1.00	2.0803×10^6	1733 s
1400	18	2.3922×10^6	0.233	1.00	56 hr	0.233	2.3885×10^6	1522 s	0.996	2.3884×10^6	1987 s

Table 2 Comparisons of Model Performances

The running times range from 2 minutes to 1.3 hours for the goal-driven model and 2 hours to 4 days (i.e., hitting our specified time limit) for the cost-concerned model. For a long-term planning problem focusing on infrastructure planning decisions possibly spanning decades, we believe that such running times are not excessive.

From Table 2, we may make several observations. First, for the goal-driven model, the optimal chance for meeting the ROI target first increases rapidly in overall flow or adoption rate, then approaches some steady level as adoption rate further increases. The objective value does not seem to be a smooth, and sometimes not even always monotonic, function of overall flow or adoption rate, despite showing an overall positive trend. When flow rates are scaled up (due to increased adoptions), there are several counteracting effects. First, the exogenous portion of variable profit increases linearly because more vehicle-miles are traveled. Second, increasing flow rates suggests that swapping demand, and thus inventory requirements at swapping stations, should go up. We may expect the inventory costs to exhibit economies of scale, due to risk-pooling effects facing random demand, allowing safety stock levels to increase at a diminishing rate as demand increases. Finally, due to the charging load capacity at stations, the service provider has to locate more stations to spread out the local charging load at individual stations. Therefore, the total fixed cost of locating stations increases, and the ROI target becomes more difficult to achieve. The interactions between these effects lead to non-smoothness of the relationship between the chance of meeting target and the adoption rate. This observation suggests that, although BP emphasizes on mass adoption as their mission, further pushing a significantly high adoption rate may not be too beneficial from a business perspective, as the service provider is required to scale up infrastructure investments to meet the extra demand.

Second, when evaluating the solution from either model using the other objective, we find that the solutions obtained from both models perform similarly well under both objectives. This suggests

that both objectives are closely correlated, under the reasonable assumptions on parameters that we make. It is also notable that sometimes the goal-driven model can even suggest solutions with better worst-case expected costs than the cost-concerned model. This is due to the fact that the algorithm for the cost-concerned model often has to be terminated (with suboptimal solutions) due to reaching the 4-day CPU time limit or the 30,000 branch-and-bound node limit. Because of this close correlation between the two objectives, we will use the goal-driven model in our subsequent tests due to the vastly-faster computation speed. One possible reason behind the significant difference in computation time between the two MISOCPs is that the cost-concerned model involves nonlinear (SOCP) constraints with binary variables, while in the goal-driven model, all nonlinear constraints involve only continuous variables. This possibly carries an impact on the performance of the solver, whose capabilities to deal with linear integer problems (and subproblems) are more sophisticated relative to integer SOCP problems (and subproblems). For example, the solver should automatically generate a richer class of valid inequalities given a set of integer linear constraints than given a set of integer second-order conic constraints (see, e.g., [2] for more discussion).

7.2. Technological Considerations

In the second experiment, we evaluate the impacts of technological advancements on the profitability, and thus economic feasibility, of the battery swapping concept. Charging operations can be greatly impacted by improvements in battery capacity (travel range), battery longevity and production costs (battery holding costs), and charging technology (charging times).

Significant improvements in battery capacity, at a reasonable price level, has been commented to be highly unlikely in the next twenty years by experts in the battery industry [15]. Therefore, we will only evaluate improvements within a small range (80 to 120 miles per charge). Note that as battery capacity increases, range extension with battery swapping is less of a concern, and therefore the service provider enjoys a drop in lower investment costs, while the exogenous portion of variable profit will not decrease. On the other hand, battery production costs are projected to drop once demand ramps up, due to economies of scale. Longevity of batteries, measured by the number of recharge cycles the battery can go through before its capacity drops to an unusable level for transportation purposes, is also projected to improve with technology. Recall that holding cost of batteries arise from both the opportunity cost of tying up capital, as well as the depreciation due to limited usage life. Reduction in production cost and improvements in longevity will decrease the cost of holding batteries for swapping, and the overall expenditure level for the service provider. Furthermore, charging times can be shortened by the use of improved chargers, before reaching a limit due to safety and longevity considerations. This can potentially be beneficial to the service provider, because as suggested by equation (2), shorter recharging times allow swapping stations to

stock fewer batteries, while attaining the same service level. These possible aspects of technological improvements carry significant impacts on the costs and the need for carrying spare batteries at swapping stations. In this experiment, we vary the values of three problem parameters: (i) the one-charge travel range (80 miles, 100 miles, 120 miles), (ii) holding cost per battery per year (\$1000, \$2000, \$3000), and (iii) charging time required (1.5, 2, 2.5, 3, 3.5 hours). The results are shown in Figure 3.

From Figure 3, we are able to draw several observations. First, profitability increases when (i) battery capacity increases, (ii) holding cost decreases and (iii) charging time decreases, as expected. Therefore, all three possible technological improvements can potentially bring significant benefits to the service provider. Second, we observe that if the charging time is long, which leads to significant battery inventory requirements, then improvements in the other two aspects do not have significant impact on profitability. In particular, in our numerical tests, the (upper bound) estimated probability of reaching a 10% annual ROI target, given by γ^* , only ranges between 0.2 to 0.4 when the charging time is 3.5 hours. To obtain a 0.5 probability guarantee, even with favorable advancements in the other two aspects, the charging time must be at most 2.5 hours. This suggests that the service provider ought to *strategically* use a high enough recharging rate. However, when charging time is already short enough (e.g., 2 hours in our tests), the potential impacts of the other two factors become significant. In practice, reductions in recharging time can be achieved by either using smaller batteries (reduction in range) or recharging at higher voltages (reduction in longevity). Our results suggest that, when recharging times are reasonably small, further reduction of recharging times, at the expense of decreasing battery longevity or capacity, is not desirable. For example, in Figure 3, when the recharging time is already shortened to 2 hours, when the battery capacity is 100 miles and the holding cost is \$1000 per year, it is not desirable to further shorten charging times (1.5 hours) at the expense of decreasing battery capacity (to 80 miles) and damaging longevity (increasing holding cost to \$2000).

7.3. Value of Standardization

Tesla, the manufacturer of the popular luxury EV “Roadster”, recently announced that its next model, the sedan “Model S”, will feature a swappable battery [27]. Yet, Tesla stressed that their philosophy is that each EV model should use its uniquely-designed battery packs, instead of standardized ones. In response, BP claimed that it will supply different battery types at its swapping stations [33]. In fact, one of the common criticisms on BP’s business model is the difficulty and high cost of accommodating different battery standards. Battery inventory requirements at swapping stations will inevitably be higher if a unified standard does not exist, because of the reduced degree of risk-pooling.

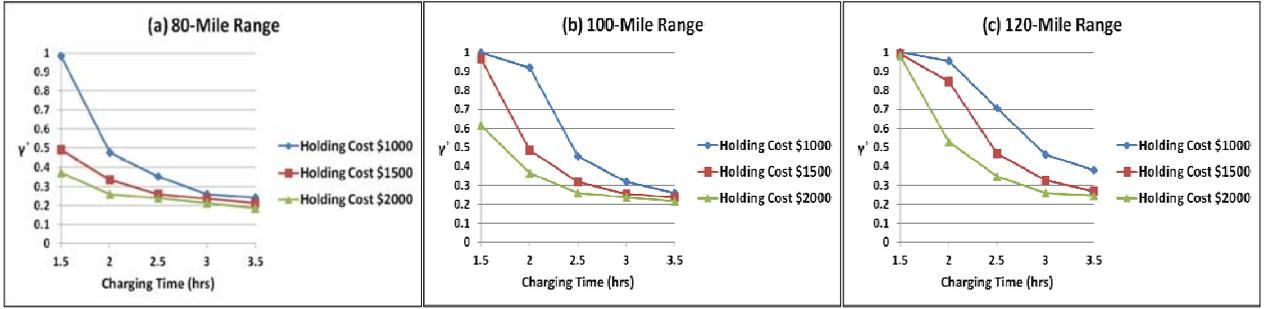


Figure 3 Impact of Technology Advancements

In this section, we consider the impact of standardization on swapping network design. If there exists a certain proportion of EVs (e.g., Tesla Model S) that use a swappable battery of a different standard from the one that the service provider uses, should the provider accommodate or ignore it? What will be the potential benefits of engaging in a common standard with the other manufacturer? In particular, we assume that a proportion $\delta \in [0, 0.5]$ of EV users require a different battery standard. Assuming that both types of batteries have the same unit holding cost h for simplicity, from (2), the total holding cost at station j becomes:

$$hI_j = t \sum_{p \in P} \lambda_p Y_{jp} + \Phi^{-1}(\alpha)(\sqrt{\delta} + \sqrt{1 - \delta}) \sqrt{t \sum_{p \in P} \lambda_p Y_{jp}}. \quad (28)$$

Because the factor $\sqrt{\delta} + \sqrt{1 - \delta} > 1$, the total holding cost increases due to the lack of a common standard. Depending on potential profits, the service provider may decide not to accommodate the alternative standard. In this computational test, we compute the values of δ at which it is profitable to accommodate both standards, and the potential profit increase if standardization is possible. For simplicity, we assume that the adoption rates of the two standards are perfectly correlated. Therefore, for a certain path p , the expected flow rates of the two types of vehicles will be $\delta\lambda_p$ and $(1 - \delta)\lambda_p$, respectively. This treatment allows us to focus on the risk-pooling effects of inventory, while avoiding complications arising from the joint variabilities regarding adoptions of the two standards.

The results are shown in Figure 4. In this experiment, we vary the profit margin per vehicle-mile of flow (this margin, multiplied by the length of the travel path, gives the exogenous marginal profit π_p), and the market share of the alternative standard (δ), while keeping the total hourly flow along the entire network constant at 600. Then, we contrast the optimal objective values of the goal-driven model (a tight upper bound on the probability of meeting ROI target, γ^*) for

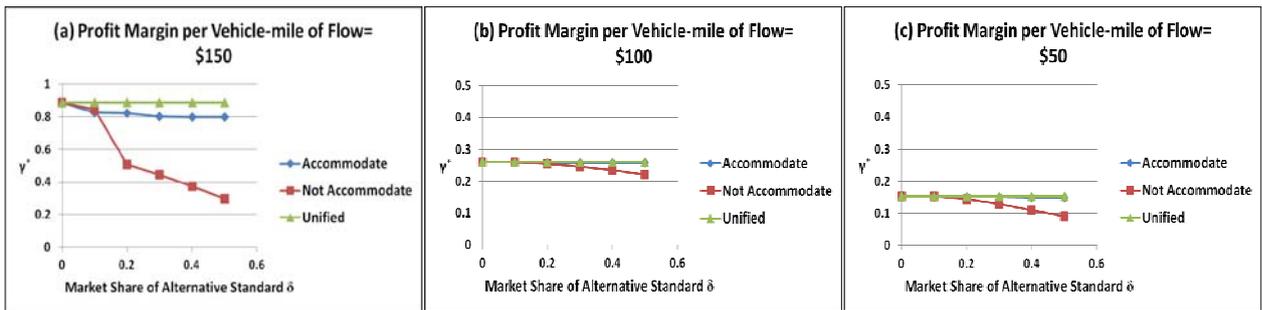


Figure 4 Impact of an Alternative Battery Standard

accommodating or not accommodating the alternative standard, as well as the case when the two standards are unified. We can draw a few observations from our results. First, profitability tends to decrease in the market share of the alternative standard δ , assuming a constant overall demand level, whether the alternative standard is accommodated or not. For the case where only the standard with larger market share is accommodated, increasing values of δ leads to lower subscription rates and lower revenue. As suggested by the results in Section 7.1, subscription rates heavily impact profitability. Therefore, the negative impact of higher δ values on profitability is significant. Note that this negative effect is more significant when the profit margin is high. For example, in Figure 4(a), the (upper bound) estimated probability of meeting the ROI target can drop from 0.9 in the unified case to 0.3 in the case of not accommodating the alternative standard. In contrast, when the profit margin is low, the (upper bound) estimated probability of meeting the ROI target is low regardless, and therefore the lack of standardization cannot do much more damage. For the case where both standards are accommodated, profitability drops as well, because the inventory risk-pooling effect diminishes as the demand for the two types of non-substitutable batteries are closer to each other. This can be easily observed from equation (28), in which the extra safety stock factor due to non-substitutability, $\sqrt{\delta} + \sqrt{1 - \delta}$, increases in δ in the range $\delta \in [0, 0.5]$.

Second, comparing the two effects mentioned above, our results suggest that the loss of revenue due to not accommodating the alternative standard tends to be more severe than the risk-pooling effect due to accommodating two battery standards. This is suggested by the observation that the profitability for not accommodating the second standard tends to fall more rapidly than that for accommodating both standards. For all levels of profit margins tested, the profitability for accommodating both standards is always higher when δ is large enough. In summary, our results suggest that the service provider should consider rejecting the alternative standard only if its market

share is very small. This is due to the fact that the battery swapping infrastructure is a cost driver, instead of a revenue driver. In particular, only a small proportion of trips by EVs carry implications on battery requirements at swapping stations, as the majority of trips are significantly shorter than the full-charge driving range [43]. When the service provider considers the trade-off between increasing subscriptions (accommodating both standards) and cutting spare battery holding costs (rejecting second standard), it should lean toward the former because rejecting the alternative standard leads to rejection of all users associated with it, while accommodating it involves increased holding cost associated with only a portion of extra demand that requires battery swapping.

8. Conclusion

In this paper, we study the problem of deploying battery swapping infrastructure for an EV charging service provider. As a major cost driver necessary for range extension, and thereby enabling EV users to enjoy a similar convenience level as they can with ICE vehicles, such an infrastructure significantly impacts the profitability of service provider. Two critical features to consider in the planning process are the spare battery inventory requirements for swapping at the stations and the uncertainty in adoption rates (i.e., demand). Under ambiguous information on demand distribution, we develop two distributionally-robust optimization models for the swapping station location problem: a cost-concerned model that minimizes the worst-case (out of all possible demand distributions) expected location and inventory costs, and a goal-driven model that maximizes the worst-case (out of all possible distributions) probability of meeting a pre-specified ROI target. Due to the battery inventory requirements, the cost (and profit) expressions include nonlinear terms in both the decision variables and primitive uncertainty terms. We develop tractable bounds that allow the approximate reformulation of the two problems as MISOCPs.

Using the two models, we perform several computational experiments with a realistic dataset created based on the San Francisco Bay Area freeway network. First, our results show that the two models generate solutions of similar qualities for *both* models. This suggests that the two objectives are correlated, at least under the reasonable assumptions that we make in the experiments. Therefore, the goal-driven model, which is more computationally-tractable, can be used for producing good solutions under both objectives. Second, our results suggest that advancements in battery (and charging) technology that speeds up recharging are critical in improving the service provider's profitability. Technological advancements that lower production cost and improve longevity and capacity of batteries only carry significant impacts on profitability if charging speed is fast enough. However, once recharging speed is improved to an acceptable level, it can be detrimental to push for further improvements at the expense of battery life or capacity. Finally, we test the impact of the existence of an alternative battery standard than the primary one offered by the service provider.

We show that, unless the market share occupied by the alternative standard is negligible, it is almost always optimal for the service provider to accommodate the alternative standard, despite the requirement to stock non-substitutable inventories of the two types of batteries at swapping stations.

This paper is the first attempt to study issues related to the interesting business model of subscription-based EV charging services. We believe that significant research opportunities for operations researchers exist in this emerging industry. Among the possible extensions of our work is the consideration of the problem in a multi-period setting. In reality, infrastructure deployment is not a one-shot decision, but rather a step-by-step roll-out process. For a brand-new service model facing uncertain demand, the ability to dynamically adjust the infrastructure plan in response to demand information collected during the roll-out process is critical. Should the service provider use an aggressive deployment strategy, such that potential customers may observe the commitment and have more confidence in adopting the service, or a more conservative strategy, with the chance of ramping up once adoption numbers grow, to minimize the risk of over-investment in case of bad response?

Our current paper focuses on the strategic decisions of *designing* the battery swapping infrastructure. One very promising research topic is on optimizing the operations of the network of charging and battery swapping facilities in real time. The swapping stations and charging spots are, in a sense, queueing systems. As recharging times are significant and demands and electric grid loads are variable during the day, it is beneficial for the service provider to smooth demand for recharging and battery swapping by directing demand to facilities with low traffic. This is possible because EVs (e.g., the ones offered by BP) are usually equipped with on-board navigation devices that can communicate with a network control center [12]. The optimal control of EV charging and battery swapping demands is a challenging, yet critical, research problem that operations researchers are in prime position to tackle.

References

- [1] Ang, M., Y.F. Lim and M. Sim (2008). A Linear Storage-Retrieval Policy for Robust Warehouse Management. Working Paper, NUS Business School, National University of Singapore.
- [2] Atamtürk, A. and V. Narayanan (2010). Conic Mixed-Integer Rounding Cuts. *Mathematical Programming, Series A*, 122, 1-20.
- [3] Bai, Y., T. Hwang, S. Kang and Y. Ouyang (2010). Biofuel Refinery Location and Supply Chain Planning under Traffic Congestion. To Appear on *Transportation Research Part B*.
- [4] Barboza, D. (August 19, 2010). China to Invest Billions in Electric and Hybrid Cars. *New York Times*. Retrieved from <http://www.nytimes.com/2010/08/20/business/energy-environment/20car.html>.

-
- [5] Ben-Tal, A. and A. Nemirovski (1998). Robust Convex Optimization. *Mathematics of Operations Research*, 23, 769-805.
- [6] Ben-Tal, A. and A. Nemirovski (1999). Robust Solutions of Uncertain Linear Programs. *Operations Research Letters*, 25, 1-13.
- [7] Berman, O., D. Bertsimas and R.C. Larson (1995). Locating Discretionary Service Facilities, II: Maximizing Market Size, Minimizing Inconvenience. *Operations Research*, 43(4), 623-632.
- [8] Berman, O., R.C. Larson and N. Fouska (1992). Optimal Location of Discretionary Service Facilities. *Transportation Science*, 26, 261-277.
- [9] Berman, O., D. Krass, C.W. Xu (1995). Locating Discretionary Service Facilities Based on Probabilistic Customer Flows. *Transportation Science*, 29(3), 276-290.
- [10] Bertsimas, D. and M. Sim (2003). Robust Discrete Optimization and Network Flows. *Mathematical Programming*, 98, 49-71.
- [11] Bertsimas, D. and M. Sim (2004). The Price of Robustness. *Operations Research*, 52(1), 35-53.
- [12] Better Place (2010). *Global Progress*, retrieved on Sep 10, 2010, from <http://www.betterplace.com/global-progress>.
- [13] Better Place (Mar 3, 2011). *Better Place and Renault Launch Fluence Z.E., the First "Unlimited Mileage" Electric Car Together with Innovative eMobility Packages, in Europe's First Better Place Center*, Press Release, Retrieved from <http://www.betterplace.com/the-company-pressroom-pressreleases-detail/index/id/better-place-and-renault-launch-fluence-z-e-the-first-unlimited-mileage-electric-car-together-with-innovative-emobility-packages-in-europe-s-first-better-place-center>.
- [14] Birge, J.R. and J.H. Dulá (1991). Bounding Separable Recourse Functions with Limited Distribution Information. *Annals of Operations Research*, 30(1), 277-298.
- [15] Boston Consulting Group (2010). Batteries for Electric Cars: Challenges, Opportunities, and the Outlook to 2020. Retrieved from <http://www.bcg.com/documents/file36615.pdf>.
- [16] Bradsher, K. (2011, Dec 26) In China, Power in Nascent Electric Car Industry. *New York Times*. Retrieved from <http://www.nytimes.com/2011/12/27/business/global/chinas-push-for-electric-cars-flows-through-grid-operators.html?pagewanted=all>.
- [17] Chambers, S. (2011, Mar 15). China Charging: Plug-in or Battery Swap? *Electric Vehicle Update*. Retrieved from <http://analysis.evupdate.com/industry-insight/china-charging-plug-or-battery-swap>.
- [18] Chen, L.R. (2007). A Design of an Optimal Battery Pulse Charge System by Frequency-Varied Technique. *IEEE Transactions on Industrial Electronics*, 54(1), 398-405.
- [19] Chen, W. and M. Sim (2009). Goal-Driven Optimization. *Operations Research*, 57(2), 342-357.

-
- [20] Chen, W., J. Sun, M. Sim and C.P. Teo (2010). From CVaR to Uncertainty Set: Implications in Joint Chance Constrained Optimization. *Operations Research*, 58, 470-485.
- [21] Chen, X., M. Sim and S. Peng (2007). A Robust Optimization Perspective on Stochastic Programming. *Operations Research*, 55(6), 1058-1071.
- [22] China Automotive Information Net (2012, Feb 1). Beijing Largest EV Charging Station Comes into Operation. Retrieved from http://english.autoinfo.gov.cn/autoinfo_eng/latest/webinfo/2012/02/1326873245539669.htm.
- [23] China Daily (2011, Mar 25). State Grid Plans to Expand Power Network in 2011. Retrieved from http://www.chinadaily.com.cn/bizchina/2011-03/25/content_12227592.htm.
- [24] China Daily (2011, Dec 15). Helping Drivers Make Switch to Electric Cars. Retrieved from http://www.chinadaily.com.cn/bizchina/2011-12/15/content_14269806.htm.
- [25] Department of Transportation, State of California (2010). *2009 Traffic Volumes on the California State Highway System*, California, <http://traffic-counts.dot.ca.gov/2009all/2009TrafficVolumes.htm>.
- [26] Electric On Wheels (2008). What Ever Happened to the EV1 Electric Car by GM? Retrieved from <http://www.electriconwheels.com/electric-car-articles.htm>.
- [27] Fehrenbacher, K. (2010, Oct 7). Where Tesla & Better Place Don't Agree. *GigaOM*, Retrieved from <http://gigaom.com/cleantech/where-tesla-better-place-dont-agree/>.
- [28] Galbraith, K. (2009, May 13). Better Place Unveils Battery Swap Station. *New York Times*, Retrieved from <http://green.blogs.nytimes.com/2009/05/13/better-place-unveils-battery-swap-station/>.
- [29] Goh, J. and M. Sim (2010). Distributionally Robust Optimization and its Tractable Approximations, *Operations Research*, 58(4), 902-917.
- [30] Green Options (2010). Introduction To Electric Cars And Electric Scooters. Retrieved from <http://www.greenoptions.com/wiki/introduction-to-electric-cars-and-electric-scooters>.
- [31] Hodgson, M.J. (1990). A Flow-Capturing Location-Allocation Model, *Geographical Analysis*, 22, 270-279.
- [32] Kang, S., H. Önal, Y. Ouyang, J. Scheffran and Ü.D. Tursun (2010). Optimizing the Biofuels Infrastructure. In M. Khanna, J. Scheffran and D. Zilberman (Eds.), *Handbook of Bioenergy Economics and Policy* (pp. 151-173). Springer.
- [33] Kerschbaum, M. (2010, Dec 7). Myth Buster: Can Our Switch Stations Accommodate Different Battery Types? *Better Place Blog*, retrieved from <http://blog.betterplace.com/2010/12/myth-buster-can-our-switch-stations-accommodate-different-battery-types/>.
- [34] Kuby, M. and S. Lim (2005). The Flow-refueling Location Problem for Alternative-Fuel Vehicles. *Socio-Economic Planning Sciences*, 39(2), 125-145.

-
- [35] Kuby, M. and S. Lim (2007). Location of Alternative-Fuel Stations Using the Flow-Refueling Location Model and Dispersion of Candidate Sites on Arcs. *Networks and Spatial Economics*, 7(2), 129-152.
- [36] Magnanti, T.L., Z.J. Shen, J. Shu, D. Simchi-Levi and C.P. Teo (2006). Inventory Placement in Acyclic Supply Chain Networks. *Operations Research Letters*, 34(2), 228-238.
- [37] MIT Electric Vehicle Team (2008). *Wheel to Well Analysis of EVs*. http://mit.edu/evt/summary_wtw.pdf.
- [38] Natarajan, K., M. Sim and J. Uichanco (2010). Tractable Robust Expected Utility and Risk Models for Portfolio Optimization. To Appear on *Mathematical Finance*.
- [39] Nemirovski, A. and A. Shapiro (2007). Convex Approximations of Chance Constrained Programs. *SIAM Journal on Optimization*, 17(4), 969-996.
- [40] Nourbakhsh, S.M. and Y. Ouyang (2010). Optimal Refueling Strategies for Locomotive Fleets in Railroad Networks. *Transportation Research Part B*, 44, 1104-1114.
- [41] Ortúzar, J. de D. and L.G. Willumsen (2001). *Modelling Transport, Third Edition*. John Wiley & Sons, New York.
- [42] People's Daily Online (2010, May 13). Public Transport to Help Fuel Spread of Charging Network. *People's Daily Online*, Retrieved from <http://english.peopledaily.com.cn/90001/90778/90860/6983257.html>.
- [43] Pearre, N.S., W. Kempton, R.L. Guensler and V.V. Elango (2011). Electric Vehicles: How Much Range is Required for a Day's Driving? *Transportation Research Part C*, In Press.
- [44] Popescu, I. (2007). Robust Mean-Covariance Solutions for Stochastic Optimization. *Operations Research*, 55(1), 98-112.
- [45] Schwartz, N.D. (2010, Jan 25). Sites to Refuel Electric Cars Gain a Big Dose of Funds. *New York Times*, Retrieved from <http://www.nytimes.com/2010/01/25/business/energy-environment/25electric.html>.
- [46] Schwartz, A. (2010, Jun 15). Better Place Demonstrated Two-Minute Electric Vehicle Battery Swap System. *Fast Company*, Retrieved from <http://www.fastcompany.com/1660290/better-place-demonstrates-two-minute-electric-vehicle-battery-swap-system-video>.
- [47] See, C.T. and M. Sim (2009). Robust Approximation to Multiperiod Inventory Management. To Appear on *Operations Research*.
- [48] K.H. Shang and J.S. Song, Newsvendor Bounds and Heuristic for Optimal Policies in Serial Supply Chains. *Management Science*, 49(5), 618-638, 2003.
- [49] Upchurch, C., M. Kuby and S. Lim (2009). A Model for Location of Capacitated Alternative-Fuel Stations. *Geographical Analysis*, 41(1), 95-106.
- [50] U.S. Environmental Protection Agency (2006). *Greenhouse Gas Emissions from the U.S. Transportation Sector: 1990-2003*. <http://www.epa.gov/otaq/climate/420r06003.pdf>.