

Electric Vehicle Charging Decisions using Only Market Trends with Persistence

Donald J. Hammerstrom
Pacific Northwest National Laboratory
donald.hammerstrom@pnnl.gov

Richard M. Pratt
Pacific Northwest National Laboratory
rmpratt@pnnl.gov

Abstract

Electric vehicle (EV) charging can take advantage of real-time electricity market price volatility. Presuming that an EV must be fully charged at a future target time, the EV should choose to charge using the lowest future electricity prices and thereby minimize electricity cost. Statistical methods must be used if forward prices are unavailable. In this case, historical prices and trends must be mined to anticipate which prices should be used to charge the EV. Price persistence, a tendency for electricity prices to inexplicably become and remain relatively high or low for extended durations, is particularly difficult to forecast and mitigate. This paper formulates and tests a pragmatic strategy for integrating conventional static statistical prices and the Bayesian propagation of price persistence from the current price to prices in the current and future hours. Simulations were conducted to test the cost effectiveness of charging strategy using real-time electricity prices.

1. Introduction

This paper addresses electric vehicle (EV) charging from the economic perspective of a vehicle owner when the expense of charging the EV is subject to dynamic, real-time electricity pricing. A charge controller exists to serve the interest of the EV owner. The charge controller has access to historical electricity prices, but it has no access to electricity load forecasts, generator costs, or transport constraints that might be available to utilities, transmission operators, market operators, or power producers.

The EV owner parks at a charging station when convenient, and the EV's owner specifies a future

time when the EV must be fully charged, which must be no earlier than the minimum time needed to fully charge the EV. The EV and charging station keep track of how much electricity must be supplied to the EV if it is to become fully charged. The charging algorithm becomes aware of a real-time price just prior to the start of a market interval, and a decision must be made whether the EV will accept the price and charge or await the unknown prices in future market intervals. If the price is accepted, the EV charges at a constant power for the duration of the market interval.

This scenario is becoming increasingly relevant as more EVs are being produced and used, moving transportation costs from fossil fuel to electric power grids. Rapid growth of renewable energy generation may increase the volatility of real-time electricity prices. Renewable generation is uncontrolled and intermittent and causes flexible resources to be dispatched out of their normal economic dispatch order. At the same time, energy storage devices like EVs, if controlled, might moderate the effects of renewable energy intermittency and thereby facilitate the continued growth of renewable energy.

While this paper is motivated by an EV charging scenario, the challenge lies primarily with one's ability to forecast near-term electricity prices. Electricity price distributions are irregular and time-variant. Even the most sophisticated electricity price forecasts are uncertain. The decision to either accept or reject an offered electricity price in a market interval is therefore dependent upon the statistical probabilities of prices in the remaining time intervals in the available charging time window. The accuracy of a point price forecast is less relevant.

Weron [1] provides a comprehensive overview of price forecasting issues and methods. Perhaps the summary is most useful toward establishing a consistent lexicon for price forecasting. The forecast approach of interest herein, according to [1], performs a short-term forecast using a hybrid reduced-form and statistical model, and is applied to real-time electricity markets. It's noted that relatively

Pacific Northwest National Laboratory is operated for the U.S. Department of Energy by Battelle Memorial Institute under Contract DEAC05-76RL01830.

few articles and conference papers address probabilistic forecasts of price densities, as addressed herein [1]. The specific form of density forecasting used herein, where future price is allocated to a few predetermined price intervals, is called threshold forecasting.

Zareipour, et al. [2] was cited by [1] as the first publication to teach threshold price forecasting. The method is claimed by [2] to be especially appropriate for demand-side management, which relies on a price threshold, not an exact price forecast. The strength of using threshold forecasting is motivated in [2] by multiple citations of unimpressive price forecasting errors, which are typically much larger than the errors of load forecasts in electric power systems. The threshold forecast's categorization accuracy was shown to compare favorably with autoregressive integrated moving average (ARIMA), transfer function, and dynamic regression models. Extraneous training inputs like system electricity demand were found to produce only modest improvements to the forecasts' categorization accuracy.

Ziel and Steinert [3] employ methods suitable for estimating probabilistic outcomes by price quantile. The approach is highly model-based, relying on rich knowledge of load and relative contributions of various renewable and conventional generation resources. The method is targeted for mid- and long-term price forecasting, but nothing seems to prevent its application to short-term price forecasting.

Jonsson et al. [4] predicts densities of day-ahead electricity prices using time-adaptive quantile regression. This method outperforms the generalist autoregressive conditional heteroscedasticity (GARCH) model. The model of [4] first uses the point prediction methods of [5] to remove seasonal and other trends; however, this underlying reliance on a point prediction necessitates fitting the residual error between the actual and predicted prices to a density distribution model—Gaussian, for example. The distribution's variance is tracked and describes the distribution centered on and symmetrical about the expected point price.

Ji et al. [6] propose a state-space-based forecasting of locational marginal prices by a system operator. The algorithm defines pricing system states, from which the impacts of important market scheduling activities may be derived. Computationally expensive Monte Carlo methods are used to estimate price transition probabilities, which may hold similarities to what is described as Bayesian propagation herein.

Conejo, Nogales, and Arroyo [7] teach a strategy for bidding conventional thermal electricity generation into a day-ahead market given uncertain,

normally distributed (or lognormal) future prices. An intermediate step, prior to bidding, is the scheduling of the resource. We do not address bidding in this paper, although the approach could be extended to bidding much as in [7]. This paper avoids any assumption of normalcy for future prices.

Luo, Huang, and Gupta [8] formulate a locational pricing practice for public EV charging services to ensure their profitability, enhance customer satisfaction, and reduce adverse impacts on the electric power grid. Customers' price elasticities are recursively learned from historical data. The authors therefore rely on real-time wholesale electricity price predictions and create multiple locational prices for EV charging service providers but defer the problem how customers are to make practical price-based charging decisions.

This paper offers these following contributions:

Upon presuming that non-utility-owned EV charging stations will not have access to influential pricing input variables (e.g., generator costs) that are available to grid operators, a short-term electricity price forecast is formulated to rely on only price trends and price persistence that may be gleaned from historical price data.

An intuitive representation of successive charging decision logic is taught and demonstrated. An EV's decision to charge or not in an impending market interval depends only on the impending interval's electricity price, forecasted price distributions, the number of market intervals remaining in the charging session, and the EV's current state of charge.

An EV charging algorithm is offered that combines advantages from both static price distributions and the propagation of transitions between defined price bins on those distributions. The static information retains price distribution irregularities; the Bayesian propagation model captures price persistence. The bins' populations dynamically grow or shrink, but the bins' price boundaries, once defined, remain constant.

The paper's method is tested using many scenarios having different initial EV states of charge, different starting times, and different charging window durations. The method is found to capture about half the theoretically achievable value in a challenging real-time market.

2. Price data characterization

Real-time, 5-minute energy prices at the Ontario node of the Independent Electricity System Operator [9] are used to demonstrate the paper's approach.

Prices from 2018 are used to train and characterize the price statistics. Then, prices from the first quarter of 2019 are used to test the strategy.

2.1. Static price correlations

We first look at static correlations. These are informed by long-term price correlations and should not include dynamic effects like lags or price persistence. We address hourly, weekly, and longer monthly or seasonal price patterns.

2.1.1. Hourly price pattern. Figure 1 shows the quartiles of 2018 prices by hour of day. A typical diurnal pattern is observed. Median prices tend to be relatively high in the morning and late afternoon. The variability of prices is also observed to be larger in hours from 6:00 to 23:00 than from 23:00 to 6:00. Remarkably high priced peak outliers can occur nearly any hour of the day.

Prices in this market frequently become negative.




Figure 1. Price quartiles by hour. A diurnal pattern is evident.

2.1.2. Weekly price pattern. Electricity consumer behavior often differs by weekday. Weekend days may have lower electricity consumption than week days. One might therefore expect to also see price variability by week day. The quartile plots of Figure 2 dispel this hypothesis. Weekday prices share very similar distributions in this market.

2.1.3. Seasonal price pattern. Figure 3 shows 2018 price quartiles by month. Monthly and seasonal price patterns are observed.




Figure 2. Price quartiles by weekday. Surprisingly, virtually no price pattern is evident by weekday.




Figure 3. Price quartiles by calendar month. Clearly there are seasonal price patterns.

2.2. Price persistence

The static distributions of the prior section would suffice if new prices were randomly drawn from these distributions. Instead, one market interval's price is highly correlated to the preceding price. Prices tend to remain relatively high or low for multiple market periods. Price persistence is represented by the prices' autocorrelation, as shown here for weeklong (Figure 4a) and daylong (Figure 4b) periods. Raw prices remain strongly correlated for several hours. The vestiges of price correlation, however, take over 4 days to fully dissipate.




Figure 4. Autocorrelation of raw price data over a) 1 week and b) 1 day.

It is precisely this persistence effect that motivates the Bayesian propagation strategy. If one is to anticipate the statistics of prices within a window of future prices so that an EV can be economically charged, both the static price distributions and models of price persistence must be used.

3. Price correlation models

A contribution of this paper is its melding of static, long-term statistical price probabilities and the management of persistence effects. The strategy put forth in this section grants moderate independence between the two statistical components.

3.1. Binned static price statistics model

The creation of correlation models using historical data is not new, and this paper claims few contributions in this respect. In fact, the simulation performance tests used later in this paper used only global hourly groups during the formulations of long-term statistical models, which may attest to the insensitivity of this strategy to detailed correlation models. We refer to these distributions as static because they need not and should not address any dynamic, intertemporal effects. Regardless how they are derived, correlation models should be constructed to capture important correlations.

This paper does, however, contribute an approach to binning of static prices that creates a foundation for the Bayesian propagation treatment of persistence, to be discussed in the next section. Ten equally spaced probability bins are overlaid on the cumulative probability price curve in Figure 5. The correct number of bins may be debatable, but 10 bins were used throughout this paper. The widths of the corresponding price bins are seen to vary greatly for this example. The feature that sets up the application of Bayesian propagation here is that the likelihoods

(i.e., forecasted populations) of the 10 price bins will be allowed to change dynamically, but the price bins do not. In a single market interval, 100% of the likelihood falls in a single bin; over long periods of time, each bin should revert to containing about 10% of the prices, the original allocation.

This practice of using static bins has an effect similar to trend removal. A trajectory from Bin 10 in one hour to Bin 10 in the following hour represents no change in persistence, but the prices in the two bins may differ, representing the typical change in prices between those two hours.

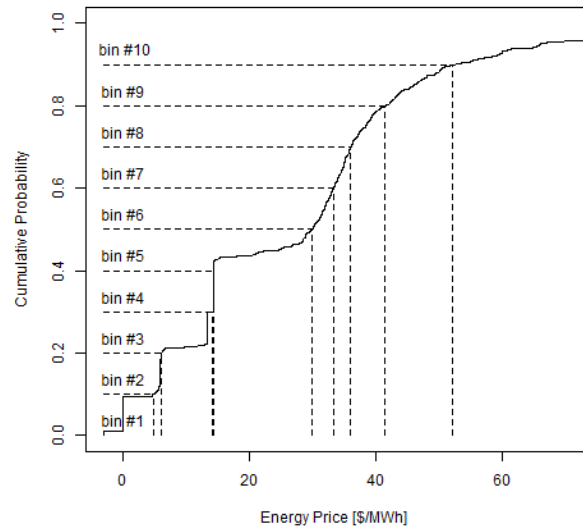


Figure 5. Hourly static bin prices. A long-term hourly distribution is divided into 10 equal-sized bins. These bins are separated by nine energy prices. The prices that demarcate the bins remain constant, but the sizes of the bins are allowed to vary due to persistence. The price statistics cause widths of price bins to vary greatly. This data happens to represent the period 10:00 – 11:00 for all 2018. The highest prices have been cropped so that price bins may be seen well.

During a training period, all the real-time, 5-minute energy prices from 2018 at the Ontario node, Independent Electricity System Operator [8] were collected. Decile price bins like those in Figure 5 were defined for each hour of the day. The method could be bootstrapped with much less historical training data. For example, the method could begin with normal distributions based on hours' price averages and standard deviations. Recursive methods could be formulated to improve the hours' actual historical price distributions over time. Forecast quality should improve as the price distributions are made to match those of the forecast hours. In this case, recent days or specific month's distributions

(referring to the variability evident in Figure 3) might have been preferable.

3.2 Bayesian propagation model

The 2018 price data was further evaluated to model the propagation of relative prices within and between successive hours. Suppose that the likelihoods of the price bins in Figure 5 are resented by a vector of probability values $P(i)$, where i is one of the ten bins. The sum of the vector members must be unity.

Given a rich history of price data, the likelihoods of transitions $P(\tilde{i}|j)$ from bin j in the first hour to \tilde{i} in another may be determined. The accent indicates bin statistics that are being propagated from other known bin probabilities.

The two hours may be the same hour, indicating the likelihood of a bin's price occurring in the hour given that another price has occurred. Often, the likelihood of an hour's price bins are statistically related instead to bin likelihoods in the prior hour. The propagation of bin probabilities may therefore be concisely stated as in (1).

$$\begin{aligned} \begin{bmatrix} P(\tilde{1}|1) & P(\tilde{1}|2) & \dots & P(\tilde{1}|10) \\ P(\tilde{2}|1) & P(\tilde{2}|2) & \dots & P(\tilde{2}|10) \\ \vdots & \vdots & \ddots & \vdots \\ P(\tilde{10}|1) & P(\tilde{10}|2) & \dots & P(\tilde{10}|10) \end{bmatrix} \begin{bmatrix} P(1) \\ P(2) \\ \vdots \\ P(10) \end{bmatrix} \\ = \begin{bmatrix} P(\tilde{1}) \\ P(\tilde{2}) \\ \vdots \\ P(\tilde{10}) \end{bmatrix} \end{aligned} \quad (1)$$

The propagation is described as *Bayesian* in that a hypothesis that a price will fall in a given bin in a future hour is updated by evidence of the recently published price and the bin in which it lies using (1). The propagation of price bin probabilities makes storage efficient. The statistics of multiple successive hours may be strung together using (1). There is no need to evaluate or store the dynamic transitions between any hours that are not contiguous.

3.3 Decision logic

At market interval $\tau = i$, an energy price λ_i becomes published and known to the EV charger. The challenge is to predict whether the published energy price in the impending market interval should be used to charge the EV. The static and dynamic,

propagated price statistics should support this decision. For each charging decision (each market interval) we calculate the probable total remaining cost of charging the EV under two cases:

- Case 1: charging occurs during the impending market interval, which incurs expected remaining cost according to (2), or
- Case 2: The EV does not charge in the impending market interval, which incurs expected remaining cost according to (3).

These two alternative costs are directly comparable. The cheapest total cost alternative should be chosen.

Each of (2) and (3) have the scaling product $\bar{p}\Delta t$, where \bar{p} is the EV charging power magnitude, and Δt is the duration of the market interval. The remaining term in parenthesis is a sum of price λ_i in the impending interval plus prices λ_τ in later intervals. The later prices are relevant only if the cumulative distribution $\Phi(\lambda_\tau)$ is less than or equal to an important ratio that represents the fraction of remaining market intervals (excluding the current impending one) that must be used to finish charging the EV. The current number of needed charging intervals n is reduced by one in the charging case (2) because the impending interval is indeed used, thus reducing the number of future market intervals needed for charging.

$$C_{\text{charge}} = \bar{p}\Delta t \left(\lambda_i + \sum_{\tau=i+1}^{\Gamma} \lambda_\tau |_{\Phi(\lambda_\tau) \leq \frac{n-1}{\Gamma-i-1}} \right) \quad (2)$$

$$C_{\text{wait}} = \bar{p}\Delta t \left(\sum_{\tau=i+1}^{\Gamma} \lambda_\tau |_{\Phi(\lambda_\tau) \leq \frac{n}{\Gamma-i-1}} \right) \quad (3)$$

4. Example charging decision scenario

This section presents details concerning a single decision whether to charge an EV during an impending 5-minute market interval for which the energy price has been announced. While it's not especially critical to the discussion, the scenario conditions happen to be the first charging interval starting at 17:20 on Feb. 3, 2019. The EV arrives at its charger 60% charged and anticipates a 1-h charging window having altogether 12 5-minute market intervals. This means that the EV must take advantage of 4.8 of the available market intervals if it is to become fully charged over the next hour.

4.1. Cumulative price distribution from component hours

The price \$24.63/MWh is determined to reside in the fourth price bin for the hour 17:00 – 18:00. The bin probabilities may therefore be found in the fourth column of the 10 x 10 matrix that represents transitions within this hour. More than 50% of the expected prices are expected to lie in price bin 4 according to this column. The likelihoods of the 10 bins are used to weigh sets of prices in the respective binned prices this hour. The bins are further weighted by the fact that 7 of the 11 of the remaining market intervals lie in the current hour. The weighted cumulative distribution for this first hour is shown in Figure 6. Its maximum contribution is $7/11=63.6\%$.




Figure 6. The Bayesian propagation method determines the weighted hourly components for all the hours in the charging opportunity window based on the relative numbers of time intervals in each hour. A cumulative price distribution is then assembled from all of its weighted hourly component distributions.

The calculation of the weighted cumulative distribution for the second hour proceeds much as for the first. However, the binned price distributions in the second hour are found by multiplying the 10 x 10 matrix of bin magnitudes in hour starting 18:00 given hour starting 17:00, and this matrix is multiplied by the vector that was calculated for the first hour prices. If the scenario duration had been longer, the process would be continued to find hourly component distributions for the additional hour or hours.

The cumulative distribution is then found by adding its hourly components. The final cumulative

distribution represents the likelihoods of prices in 11 available market intervals excluding the current one for which the price has already been published.

4.2. Example charging decision

The cumulative price distribution for the available future charging time intervals may now support decisions to either charge during the current time interval or wait for anticipated cheaper prices. Figure 7 compares the respective calculations of waiting cost (3) and charging cost (2). If the EV were to wait, it must still use 4.8 of the remaining 11 charging intervals, or 44%. If, however, the EV were to charge in the impending time interval, it would incur the impending interval's cost, and need only 3.8 of the remaining 11 market intervals, or 35%. Therefore, the decision to charge or not has implications for the average statistical costs of the intervals that must still be used. The total costs are $\$15.89 \cdot p$ if charging² and $\$15.95 \cdot p$ if waiting.³ The decision is close, but statistically speaking, the EV should charge in the impending market interval.

5. Simulation of 3600 test scenarios

Simulation scripts were coded in R statistical programming language [9]. This section discusses the simulation scenarios that were set up and run to compare the performances of this paper's EV charging strategy with two alternative strategies.

5.1. Methods to be compared

Three alternative EV charging strategies are compared:

- This paper's Bayesian propagation strategy,
- A conventional, immediate charging strategy, in which the first available market intervals and prices are always used until the EV has been fully charged, and
- An ideal charging strategy, in which the cheapest prices can be perfectly predicted for any future charging time window.

5.2. Scenario design

Five-minute, real-time price data was obtained from the Independent Electricity System Operator

² $(p \text{ [MW]})(0.0833 \text{ [h]})(\$24.63 \text{ [$/MWh]} + 11 * 15.1 \text{ [$/MWh]})$
³ $(p \text{ [MW]})(0.0833 \text{ [h]})(11 * 17.4 \text{ [$/MWh]})$

website [8], using Ontario node energy prices and sample days from the first quarter of 2019. A set of 3600 test scenarios was formulated from permutations of four categorical variable sets that are detailed in Table 1.

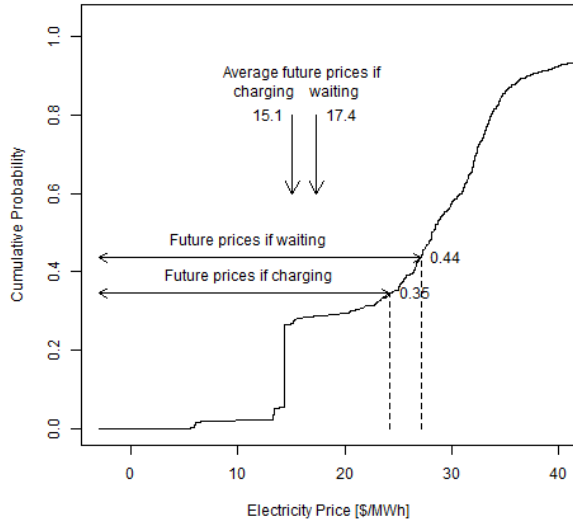


Figure 7. Example charge / no charge decision. In this example, a decision to charge means that only 35% of the future time intervals must be used to complete the charge, while waiting means that 44% of them must be utilized. The decision affects the anticipated set of low remaining prices that must be used to charge the EV. The areas under the charging and waiting cumulative distribution alternatives yield average anticipated prices for the charging and waiting alternatives. Finally, one may calculate the sum predicted cost of charging, including the cost of using the current market interval, with that of waiting.

Table 1. Categorical scenario variables

VARIABLE	#	SET
Day	30	Jan. 1, Jan. 4, Jan. 7, ... Mar. 31, 2019
Duration	5	1, 2, 4, 8, and 16 [h]
State of charge	4	20, 40, 60 and 80 [%]
Start time	6	0:00, 4:20, 8:40, 13:00, 17:20, and 21:40

5.3. Test days and test day prices

A set of 30 scenario days were selected using market prices from every third day of the first quarter of 2019. The quartile plot of Figure 8 confirms how truly diverse and volatile the prices were during the selected test days.




Figure 8. Statistics of all energy prices for the 30 selected scenario days. Prices are shown to be diverse, which makes the predictions of prices during EV charging opportunity windows particularly challenging.

5.4. Impact of time that charging begins

The scenarios were run at a variety of different times during the day to make sure that the impacts of diurnal cycles could be tested. Figure 9 compares quartile plots for the average charging prices of the three EV charging strategies at the tested starting times. The prices reflect the normal diurnal price trend as in Figure 1. The performances of the alternative strategies are consistently ordered from conventional, Bayesian, and ideal strategies. This paper's Bayesian propagation strategy median prices often lie approximately half way between those of the other two.

5.5. Impact of charging opportunity window duration

The three charging strategies were compared in respect to the duration of the charging opportunity window duration. The luxury of having a long time before the EV must be charged means that one might be pickier about which charging prices to accept. This was indeed the case for both the Bayesian propagation and ideal charging strategies in Figure 10. As should be expected, the immediate charging strategy did not benefit from having longer charging opportunity windows. The vehicle using immediate charging would simply become fully charged until the end of an available charging window.




Figure 9. Quartile plot of average EV charging prices for the three charging strategies as a function of the time of day when charging commences.




Figure 10. Quartile plots of average charging price for the three EV charging strategies as a function of the duration of the charging opportunity window.

5.6. Impact of initial EV state of charge

The tested performances of the three charging strategies were compared in respect to vehicles' states of charge at the initiation of the charging window. A vehicle that is nearly fully charged at the start of its charging scenario can be relatively selective and wait for the least expensive electricity prices. The EV that arrives nearly depleted, however, must use many of the charging intervals—even if prices are high. This effect is strongly evident in Figure 11 for the ideal charging strategy, less so for

the paper's Bayesian propagation strategy, and not at all for the immediate charging strategy.




Figure 11. Quartiles of the average scenario charging prices for the three charging strategies as a function of the EV's initial states of charge.

5.7. Comparison of strategies

Figure 12 directly compares the average charging prices of the paper's charging strategy and the conventional charging strategy. The paper's strategy tended to pick cheaper electricity prices, on average, than did conventional strategy, but not always.

The degree to which the paper's strategy fell short of the ideal is shown in Figure 13. Many of its scenario prices lie close to those of the ideal, but outliers reveal that the paper's strategy is fallible and cannot always predict the statistics of prices in the charging opportunity window.

Figure 14 compares the average charging prices using the paper's strategy with the average price of charging intervals that have been chosen randomly from those in the available charging time window. This figure looks different each time it is run because of the random selection. The paper's selection strategy is shown to be superior to a random strategy for most of the simulated charging periods.

Based on summed charging prices for the entire set of 3600 test scenarios, the paper's strategy reduced the cost of charging in this market by 11.9%, on average, over that of conventional charging. Interestingly, the paper's strategy cost 12.0% more than the ideal. It is reasonable to say, therefore, that the paper's strategy recaptures about half the

theoretical potential economic benefit that is missed by the simple conventional charging strategy.




Figure 12. Normalized average electricity cost for the Bayesian propagation process versus conventional EV charging using the first available market prices. The data includes 3600 EV charging scenarios over 30 days selected from the first three months of 2019.

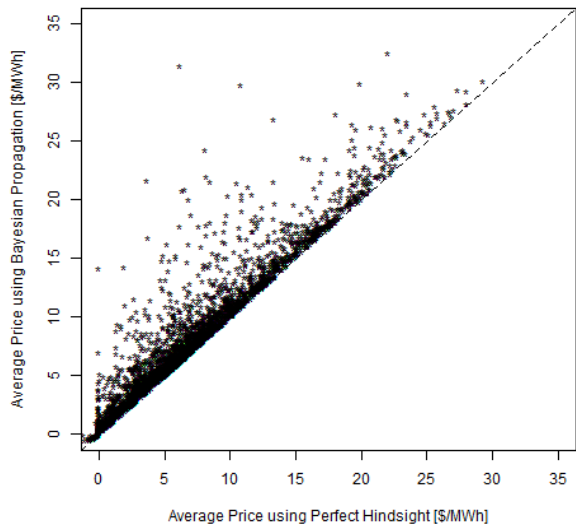


Figure 13. Average charging prices using the Bayesian propagation method versus cheapest prices chosen with prescient foresight.




Figure 14. Average charging prices using the Bayesian propagation method versus when intervals are randomly selected.

6. Conclusion

This paper addressed the challenge of selecting the most economical dynamic electricity market prices to charge an EV, given that an EV arrives, monitors its state of charge, and can state the time at which the EV must be fully charged. A statistical charging strategy is described and takes advantage of both long-term price statistics and persistence effects. The long-term statistics are addressed and trained using conventional methods, but persistence effects are modeled using the Bayesian propagation of the likelihoods of transitions between current and successive hourly price bins. An entire year of market data was available to calibrate the model. Then, the strategy was tested using new market prices. A set of 3600 charging scenarios was defined, and the performance of the paper's charging strategy was compared with the performances of the immediate charging strategy and an ideal strategy having perfect foresight. The paper's strategy was found to capture, on average, about half the theoretical economic benefits that would be lost using the immediate charging strategy. The new strategy's selected charging prices were, on average, almost 12% cheaper than those used by the immediate charging strategy in this particular market.

7. References

- [1] R. Weron, "Electricity price forecasting: A review of the state-of-the-art with a look into the future," *International Journal of Forecasting*, vol. 30, pp. 1030 - 1081, 2014.
- [2] H. Zareipour, A. Janjani, H. Leung, A. Motamedi and A. Schellenberg, "Classification of future electricity market prices," *IEEE Transactions on Power Systems*, vol. 26, no. 1, pp. 165 - 173, February 2011.
- [3] F. Ziel and R. Steinert, "Probabilistic mid- and long-term electricity price forecasting," *Renewable and Sustainable Energy Reviews*, vol. 94, pp. 251 - 266, October 2018.
- [4] T. Jónsson, P. Pinson, H. Madsen and H. A. Nielsen, "Predictive densities for day-ahead electricity prices using time-adaptive quantile regression," *Energies*, vol. 7, no. 9, pp. 5523-5547, 2014.
- [5] T. Jónsson, P. Pinson, H. A. Nielsen, H. Madsen and T. S. Nielsen, "Forecasting electricity spot prices accounting for wind power predictions," *IEEE Trans. Sustain. Energy*, vol. 4, pp. 210-218, 2013.
- [6] Y. Ji, J. Kim, R. J. Thomas and L. Tong, "Forecasting real-time locational marginal price: A state space approach," in *2013 Asilomar Conference on Signals, Systems and Computers*, 2013.
- [7] A. J. Conejo, F. J. Nogales and J. M. Arroyo, "Price-taker bidding strategy under price uncertainty," *IEEE Transactions on Power Systems*, vol. 17, no. 4, pp. 1081 - 1088, 2002.
- [8] C. Luo, H. Yih-Fang and G. Vijay, "Stochastic dynamic pricing for EV charging stations with renewable energy integration and energy storage," in *IEEE Transactions on Smart Grid*, 2017.
- [9] Independent Electricity System Operator, "Yearly real time market price report," 18 April 2019. [Online]. Available: <http://reports.ieso.ca/public/RealtimeMktPriceYear/>. [Accessed 18 April 2019].
- [10] R Core Team, "R: A language and environment for statistical computing," 2018. [Online]. Available: <https://www.R-project.org/>.
- [11] "Volttron," [Online].
- [12] R. J. Hamidi and H. Livani, "Myopic real-time decentralized charging management of plug-in hybrid electric vehicles," *Electric Power Systems Research*, vol. 143, pp. 522-532, 2017.
- [13] F. Ziel, R. Steinert and S. Husmann, "Efficient modeling and forecasting of electricity spot prices," *Energy Economics*, vol. 47, pp. 98 - 111, 2015.