Risk-Sensitive Peak-Aware Energy Scheduling: Competitive and Learning-Augmented Algorithms

Lukas Himmelreich
ETH Zurich
Zurich, Switzerland
lukas.himmelreich@inf.ethz.ch

Nicolas Christianson
Stanford University
Stanford, California, USA
christianson@stanford.edu

Adam Wierman
California Institute of Technology
Pasadena, California, USA
adamw@caltech.edu

ABSTRACT

Risk is a critical consideration in many energy-related problems, where stakeholders are often sensitive to rare but costly events. However, traditional online algorithms typically optimize expected cost and can perform poorly in such risk-sensitive settings. Motivated by this challenge, we study the design of algorithms for peak-aware energy scheduling, where a microgrid operator must decide how much electricity to generate locally vs. purchase from the grid while facing both a spot price and peak charge for grid electricity. Modeling this as an instance of the online Bahncard problem, we obtain three new results: (1) a family of closed-form algorithms with a provable bound on the Conditional Value-at-Risk (CVaR)-competitive ratio, a recently proposed risk-sensitive performance metric for online algorithms; (2) an *optimal* online algorithm resulting from the solution to a certain delay differential equation, which can be approximated numerically on an instance-by-instance basis; and (3) a *learning-augmented* algorithm for this problem which, given a hyperparameter $\lambda \in (0, 1]$, achieves $1+\Theta(\lambda)$-consistency while maintaining a risk-sensitive CVaR-robustness of $\Theta(1/\lambda)$.

1. INTRODUCTION

Online decision-making under uncertainty is central to modern energy systems. A common challenge in this area is the trade-off between repeated decisions or payments at a known (small) cost and a one-time (large) investment that yields future discounts, such as in peak-aware energy scheduling problems [17]. This fundamental structure is captured by the online Bahncard problem, a classic online problem in which one must decide whether to pay full price for train tickets or to invest in the Bahncard, which gives a discount for purchasing tickets [11].

A key technique for improving performance in such online problems is randomization, which enables improved bounds in the Bahncard problem and many other online problems [11, 13, 9, 4, 5]. However, these improvements are achieved only in expectation, while performance can vary significantly in individual runs; this variability can be problematic for risk-sensitive decision-makers who are sensitive to unlikely yet very costly outcomes. To address this challenge, several recent works have investigated the design of risk-aware online algorithms [10, 8, 7, 6]. In particular, Christianson et al. [6] propose a risk-sensitive competitive ratio metric based on the Conditional Value-at-Risk (CVaR) [16], which replaces expected cost with tail cost.

To this point, the design of risk-sensitive online algorithms has been largely limited to simple online problems like ski rental and one-max search. Moreover, almost no work has integrated risk sensitivity into the design of *learning-augmented algorithms* [15, 14]; the only exception is the recent work of [2], which considers risk sensitivity with respect to the randomness of a distributional prediction. Thus, the goal of our work is twofold: first, to *push the understanding of risk-sensitive online algorithm design into the more general setting of the Bahncard problem and its applications to peak-aware energy scheduling*; and second, to *propose a notion of risk-sensitive robustness for the design of learning-augmented algorithms*.

This extended abstract gives a brief summary of our results on risk-sensitive competitive and learning-augmented algorithms for the Bahncard problem. For a full overview of our results, their application to peak-aware energy scheduling, proofs, and experiments, see the full paper [12].

Learning-Augmented Algorithms: Theory and Applications (LATA) 2026
Ann Arbor, MI, USA

2. BACKGROUND

We briefly introduce the performance metrics and problem studied in our work. To begin, we define the conditional value-at-risk (CVaR) [3, 16, 1].

DEFINITION 2.1. *Let X be a real-valued random variable with CDF F_X. For a risk sensitivity level $\delta \in [0, 1)$, the CVaR_δ is the expected value of X given that it falls within the worst $(1 - \delta)$-fraction of outcomes:*

$$\mathrm{CVaR}_\delta[X] = \frac{1}{1-\delta} \int_\delta^1 F_X^{-1}(p)\,\mathrm{d}p. \tag{1}$$

We now define the risk-sensitive performance metric we focus on: the CVaR-competitive ratio [6].

DEFINITION 2.2. *The CVaR_δ-Competitive Ratio (δ-CR) of a randomized online algorithm* ALG *measures the ratio between the CVaR_δ of its cost and the optimal offline solution. It is defined as:*

$$\delta\text{-CR}(\textsc{Alg}) := \sup_{\sigma\in\Sigma} \frac{\mathrm{CVaR}_\delta[\textsc{Alg}(\sigma)]}{\textsc{Opt}(\sigma)},$$

where $\mathrm{CVaR}_\delta[\textsc{Alg}(\sigma)]$ denotes the CVaR of the cost incurred by ALG *on input σ, and $\textsc{Opt}(\sigma)$ is the cost of the optimal offline algorithm on the same input.*

In the *learning-augmented* setting, our algorithm has access to prediction(s) about the problem instance, which may be used to try to improve performance. In this setting, the typical performance metrics are consistency—performance when the prediction is perfect—and robustness—performance under arbitrarily poor predictions. In this work, we employ the standard (expected cost) notion of consistency, and a risk-sensitive notion of robustness.

Definition 2.3. *A learning-augmented algorithm* Alg *is c-consistent if, when the prediction is perfect, we have*

$$\mathbb{E}[\textsc{Alg}(\sigma)]/\textsc{Opt}(\sigma) \le c.$$

Definition 2.4. *Let* Alg *be a learning-augmented algorithm and* $\delta \in [0, 1)$. *The* CVaR_δ*-robustness (*δ-ROB*) is the worst-case* δ-CR *achieved by* Alg *under arbitrary prediction error:*

$$\delta\text{-ROB}(\textsc{Alg}) := \sup_{\sigma\in\Sigma} \frac{\mathrm{CVaR}_\delta[\textsc{Alg}(\sigma)]}{\textsc{Opt}(\sigma)}.$$

Finally, we introduce the main problem we study: the continuous-time Bahncard problem with permanent discount [11]. Let $C > 0$ and $\beta \in [0, 1)$. There is some (*a priori* unknown) problem duration $T \in \mathbb{R}_+$; the decision-maker pays \$1 per unit of time until optionally purchasing a Bahncard for cost $\$C$, after which the cost per unit of time is $\$\beta < 1$. The objective is to decide when, if ever, to buy the Bahncard so as to minimize the competitive ratio. A randomized algorithm for this problem can be characterized by a random variable X governing the time at which the Bahncard is purchased. In the learning-augmented setting, the prediction will be an estimate of the normalized threshold $\tau_T := \frac{T(1-\beta)}{C}$; in the following, we will refer to "normalized times" accordingly. Note that $\beta = 0$ corresponds to the continuous-time ski rental setting studied in [6], so we will restrict our focus to $\beta \in (0, 1)$.

3. RESULTS

Our first main result characterizes the optimal algorithm whose Bahncard purchase is limited to normalized times in $[0, 1] \cup \{\infty\}$.

Theorem 3.1. *Fix* $\delta \in [0, 1)$*, and let* $\alpha_\delta^{[0,1]\cup\infty}$ *be the optimal* δ-CR *of any continuous strategy supported on* $[0, 1] \cup \{\infty\}$*, with* p_∞ *the corresponding optimal probability mass at* ∞*. Let* $\phi : [0, 1 - p_\infty] \to [0, 1]$ *be the solution to the delay differential equation defined as:*

$$\phi'(z) = \frac{1-\beta}{(1-\delta)(\alpha-\beta)}\left[\phi(z) - \phi(z-(1-\delta))\right] \quad \text{for } z > 1-\delta,$$

with initial condition $\phi(z) = \log\left(1 + \frac{z(1-\beta)}{(\alpha-1)(1-\delta)}\right)$ *on* $z \le 1 - \delta$*. Then, when* $\alpha = \alpha_\delta^{[0,1]\cup\infty}$*,* ϕ *is exactly the inverse CDF of the optimal continuous algorithm with support on* $[0, 1] \cup \{\infty\}$ *for the Bahncard problem.*

Since the above theorem does not give an analytical expression for the δ-CR, we next propose an algorithm family for which the δ-CR can be computed in closed form.

Theorem 3.2. *Let* $\delta \in [0, 1)$*. Define a probability distribution* p_δ *on* $[0, 1] \cup \{\infty\}$ *with density*

$$p_\delta(x) = \frac{\left(1 - e^{-\frac{c}{1-\delta}}\right)(1-p_\infty)(1-\delta)}{c\left(1 + \left(e^{-\frac{c}{1-\delta}} - 1\right)x\right)}, \quad \text{for } x \in [0, 1]$$

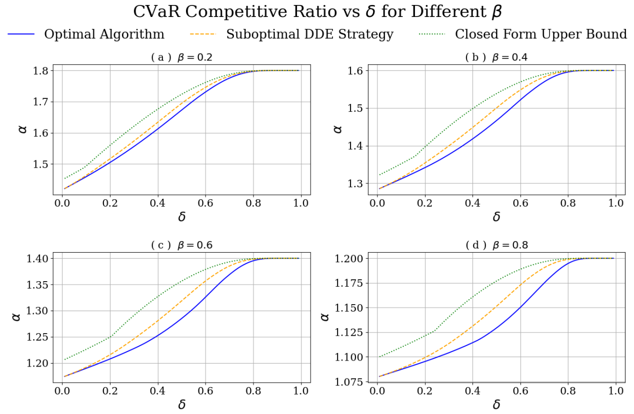

Figure 1: CVaR_δ-competitive ratio from suboptimal delay differential approach in Theorem 3.1, suboptimal closed form upper bound from Theorem 3.2, and optimal algorithm from Theorem 3.3 approximated with bisection search.

and a point mass $p_\infty = \frac{\beta(1-\delta)}{e-1+\beta}$ *at* ∞*, where* $c \approx 1.25643$*. The algorithm following* p_δ *has* δ-CR *bounded by*

$$\alpha_\delta^{p_\delta} = \begin{cases} 1 + (1-\beta)\left(\frac{2-e^{-\frac{c}{1-\delta}}}{1-e^{-D}} - 1\right) - \frac{1-\beta}{D} & \text{if } \delta \le p_\infty, \\ \beta + \frac{(1-\beta)(2-e^{-\frac{c}{1-\delta}})}{1-e^{-\frac{c}{1-\delta}}} + \frac{(1-\beta)e^{-\frac{c}{1-\delta}}(1-e^{D})}{(1-e^{-\frac{c}{1-\delta}})D} & \text{otherwise,} \end{cases}$$

where $D := c\frac{e-1+\beta}{e-1+\beta\cdot\delta}$*.*

By relaxing the support restriction to $[0, 1] \cup \{\infty\}$, we can obtain a significantly better algorithm, defined implicitly in terms of a delay differential equation.

Theorem 3.3. *For any* $\delta \in [0, 1)$*, let* $F_X^\alpha : [0, \infty] \to [0, 1]$ *be the solution to the delay differential equation in [12, Theorem 4.1]. If* α *is equal to the optimal* δ-CR *and if the corresponding solution* F_X^α *is a valid CDF, then* F_X^α *is an optimal continuous algorithm for the Bahncard problem.*

Note that we do not prove existence or uniqueness of an α yielding a valid CDF, due to the challenging structure of the delay differential equation in the previous theorem. However, in practice it appears that the optimal α can be estimated via bisection search, and the resulting CDF is valid. Figure 1 compares the CVaR_δ-competitive ratios achieved by the three strategies described thus far.

We conclude by stating our learning-augmented bound; due to space, we omit the full algorithm description.

Theorem 3.4. *For* $\lambda \in (0, 1]$*, there is a learning-augmented algorithm [12, Algorithm 1] that achieves* $1+\Theta(\lambda)$*-consistency and* $\Theta\left(\frac{1}{\lambda}\right)$*-*$\mathrm{CVaR}_\delta$*-robustness.*

In the full paper [12], we evaluate this learning-augmented algorithm in a case study on peak-aware energy scheduling using real-world datacenter and energy traces, comparing against a prior non-risk-sensitive learning-augmented algorithm and the algorithm from Theorem 3.3. Overall, our strategy exhibits good performance and improved tail behavior when compared to the non-risk-sensitive method. Interesting avenues for future work include (a) integrating uncertainty quantification into risk-sensitive algorithm design, and (b) risk-sensitive, learning-augmented algorithm design for other problems.

4. REFERENCES

[1] Carlo Acerbi and Dirk Tasche. On the coherence of expected shortfall. *Journal of Banking & Finance*, 26(7):1487–1503, July 2002.

[2] Spyros Angelopoulos, Christoph Dürr, and Georgii Melidi. Decision-theoretic approaches for improved learning-augmented algorithms. In *The Fourteenth International Conference on Learning Representations*, 2026.

[3] Philippe Artzner, Freddy Delbaen, Jean-Marc Eber, and David Heath. Coherent Measures of Risk. *Mathematical Finance*, 9(3):203–228, July 1999.

[4] Allan Borodin, Nathan Linial, and Michael E. Saks. An optimal on-line algorithm for metrical task system. *Journal of the ACM*, 39(4):745–763, October 1992.

[5] Sébastien Bubeck, Michael B. Cohen, James R. Lee, and Yin Tat Lee. Metrical Task Systems on Trees via Mirror Descent and Unfair Gluing. *SIAM Journal on Computing*, 50(3):909–923, January 2021.

[6] Nicolas Christianson, Bo Sun, Steven Low, and Adam Wierman. Risk-sensitive online algorithms, 2024.

[7] Qiming Cui and Michael Dinitz. Controlling tail risk in two-slope ski rental, 2025.

[8] Michael Dinitz, Sungjin Im, Thomas Lavastida, Benjamin Moseley, and Sergei Vassilvitskii. Controlling Tail Risk in Online Ski-Rental. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 4247–4263, 2024.

[9] R. El-Yaniv, A. Fiat, R. M. Karp, and G. Turpin. Optimal Search and One-Way Trading Online Algorithms. *Algorithmica*, 30(1):101–139, May 2001.

[10] Eyal Even-Dar, Michael Kearns, and Jennifer Wortman. Risk-sensitive online learning. In José L. Balcázar, Philip M. Long, and Frank Stephan, editors, *Algorithmic Learning Theory*, pages 199–213, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.

[11] Rudolf Fleischer. On the bahncard problem. *Theoretical Computer Science*, 268(1):161–174, 2001. On-line Algorithms '98.

[12] Lukas Himmelreich, Nicolas Christianson, and Adam Wierman. Risk-sensitive peak-aware energy scheduling: Competitive and learning-augmented algorithms. In *Proceedings of the Seventeenth ACM International Conference on Future and Sustainable Energy Systems*, e-Energy '26. ACM, 2026, to appear.

[13] A. R. Karlin, M. S. Manasse, L. A. McGeoch, and S. Owicki. Competitive randomized algorithms for nonuniform problems. *Algorithmica*, 11(6):542–571, June 1994.

[14] Thodoris Lykouris and Sergei Vassilvitskii. Competitive Caching with Machine Learned Advice. In *Proceedings of the 35th International Conference on Machine Learning*, pages 3296–3305. PMLR, July 2018.

[15] Manish Purohit, Zoya Svitkina, and Ravi Kumar. Improving Online Algorithms via ML Predictions. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.

[16] R. Tyrrell Rockafellar and Stanislav Uryasev. Optimization of conditional value-at-risk. *The Journal of Risk*, 2(3):21–41, 2000.

[17] Ying Zhang, Mohammad H. Hajiesmaili, Sinan Cai, Minghua Chen, and Qi Zhu. Peak-aware online economic dispatching for microgrids. *IEEE Transactions on Smart Grid*, 9(1):323–335, 2018.