into account. Therefore, the 1.44 percentage point increase from *EM-7* to *BADP-w* increases the net profit much more and may even turn a loss into a profit. Because of the high investment and running costs of energy storage systems they usually have a low internal rate of return (IRR) of 0.2% to 9.1% (Steffen (2012), Zhu et al. (2013), Olaszi & Ladanyi (2017)). The actual IRR depends on the specifications and costs of the energy storage and the trading strategy. If one can increase the "contribution margin" with *BADP-w* by 1.44 percentage points this also increases the IRR by nearly 1.44 percentage points. Therefore, the IRR can be dramatically increased by 16 to 840 %.

6.3. Decision making insights

To investigate the decisions of the storage owner and to show that the decisions derived by *BADP-w* are reasonable, we analyze the storage level for *Autumn*. Since we derived 20 different policies with our best approach *BADP-w* we report the average storage level over the 20 policies in Figure 6. Additionally, we report the maximum and minimum storage levels as well as the day-ahead prices. As reference points the Mondays (0:00) of *Autumn* are marked.

Figure 6: Storage level over time with prices

Overall, we observe strong daily and weekly patterns in the storage level, which follow electricity prices. On a weekly level, the storage is charged on weekends and discharged at the beginning of the week. On a daily level, it is charged right after midnight and discharged during the day. In particular, we see that at the first three Sundays the storage is full, as it was charged during the weekends where the prices are usually low to sell electricity during the week where the prices are higher. Only during the fourth weekend the storage is not fully charged as prices did not drop.

After the storage is fully charged, not all energy is sold directly on Monday. This is only true for the first Monday. After the second Monday, most energy is sold Tuesday and after the third Monday at Wednesday. These were the weekdays where the prices peaked.

In some hours or even days the gap between the policies with maximum storage level and minimum storage level is quite large (e.g. in the week of the second Monday). This happens if multiple nearly equally good decisions exist (e.g. discharging a day later at nearly equal prices).

6.4. Influence of the price forecast

To check if it is worth to handle a high-dimensional price forecast in the dynamic program, we vary the number of days in the price forecast and therefore the dimension of the exogenous part of the state $P_t^{history}$. We denote the number of days in the price history as D. The corresponding price history is $P_t^{history} = \left(P_{t+1-D}^{DA}, \cdots, P_t^{DA}, P_{t+1-D}^{ID}, \cdots, P_t^{ID}\right)$. For each $D \in \{1, \cdots, 10\}$ we estimate the price forecast using lasso with ten-fold cross validation 20 times and learn the value functions with our best approach *BADP-w*. Again we report the mean profit and the 95 % confidence bands. In addition, we use *EM-7* with seven days in the price history as a benchmark, since this does not suffer from the curses of dimensionality.

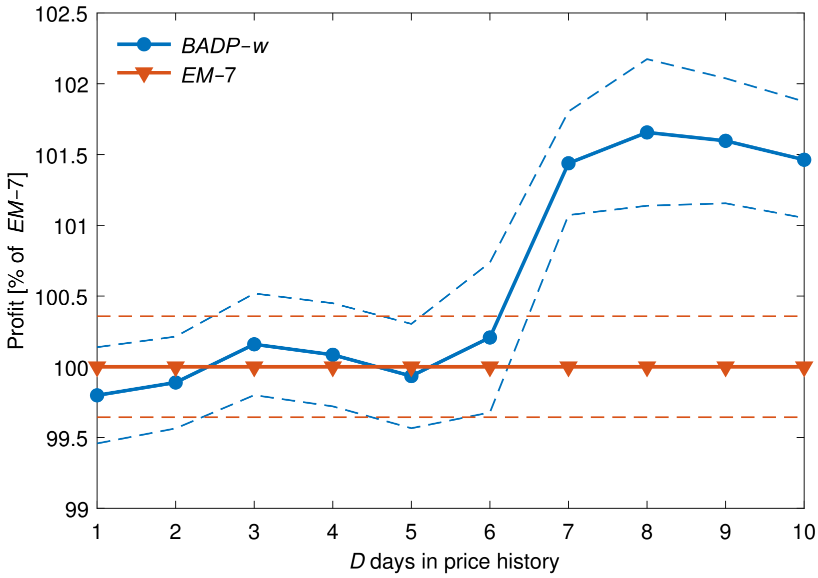

Figure 7: Influence of price forecast used in optimization on expected profits

Figure 7 illustrates the influence of the length of the price history used in the optimization on profit. We estimate ten different price processes for each market with one ($S_t^{DA} \in \mathbb{R}^{122}$) up to ten days ($S_t^{DA} \in \mathbb{R}^{1202}$) in the price history. We fit the value functions with *BADP-w* and derive the decisions based on this forecast and evaluate these with the data in our test data set. The difference of the mean profits with 1 to 6 days in the price history is not significant. But the seventh day in the price history is important for *BADP-w* to outperform *EM-7* which always uses seven days in the price history to calculate the forecast used in the expectation model. This significant increase in the profit follows because the electricity prices have a strong weekly pattern and, thus, the prices of the same day of the last week are good predictors for current prices. Again, the difference of the mean profits with 7 to 10 days in the price history is not significant. Note that the fact that profit increases in the price history is not obvious, because a larger state variable is harder to handle, which could also lead to a decline in profit. Thus, Figure 7 demonstrates that *BADP-w* can handle a large state variable. More importantly, as there is no reason not to use a high-dimensional price forecast for the benchmark approach, Figure 7 indicates that the dynamic program needs at least seven days in the price history to outperform the benchmark approach with seven days. Thus, if we are not willing to handle the large state variable in the dynamic program, the much simpler benchmark approach would be the better choice. Moreover, we also tested the influence of the price forecast on *EM-7* (i.e. with D days in the history used to compute the expectations for the next 7 days); not shown here. The profit of *EM-7* shows a similar pattern as the profit of *BADP-w*, but *BADP-w* outperforms *EM-7* for every D.

Since Figure 7 indicates that the seventh day in the price history is more important than the previous days, this suggests the reduction of the state space via feature selection. One could only include the most relevant prices of the previous week. While the importance of features (dimensions of the price history) is considered in the weights of *BADP-w*, the feature selection is not investigated in full detail. As the modeled price processes must follow the Markov property, the feature selection is not straightforward. It is possible to fit the value functions with *BADP-w* based on a simplified forecast and therefore simplified model (e.g. $D = 1$) but derive the decisions based on a more complex forecast (e.g. $D = 7$). For this, the higher dimensional S_t^{DA} and therefore P_t^{histoy} is treated as a latent state variable while deriving the trading decisions based on the fixed value function approximations. Therefore, we compare this approach with the data in Figure 7. For this, we derive the decisions based on our basic forecast ($D = 7$).

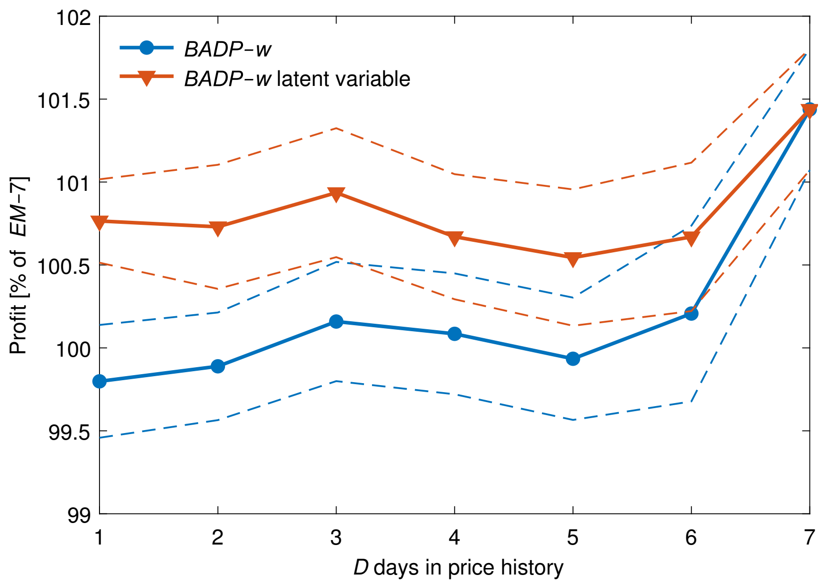

Figure 8: Influence of underlying price process on expected profits

Figure 8 is similar to Figure 7, but is extended by the latent state variable approach where the decisions for $D < 7$ are made based on the $D = 7$ forecast (basic setting). To keep it simple, we do not show the *EM-7* benchmark with its confidence bands. Figure 8 demonstrates that treating a higher dimensional forecast as a latent variable pays off, but the gain decreases in D. The latent state variable approach with three days in the price history comes closest to the basic setting $D = 7$, but further tests (not shown here) indicate that the difference is still significant. We do not include $D > 7$, because in this cases the price history would contain dimensions that are not used for the forecast and therefore would be the irrelevant.

In conclusion, it is essential for dynamic programs to use a competitive forecast to outperform a simple benchmark approach.

6.5. Gains from integrated trading

In this section, we compare the integrated day-ahead and intraday trading with restricted trading settings. In particular we consider the following settings:

- **Day-ahead only (DA):** This setting trades only in the day-ahead market and corresponds to our policy functions $X_t^{DA,\pi}(\cdot)$ and $X_t^{ID,\pi}(\cdot)$ with additional constraint $x_{t,h,q}^{ID} = 0\ \forall h, q$.
- **Intraday only (ID):** This setting trades only in the intraday market. This can be reflected with our policy functions $X_t^{DA,\pi}(\cdot)$ and $X_t^{ID,\pi}(\cdot)$ by adding the constraint $x_{t,h}^{DA} = 0\ \forall h$.
- **DA+ID, sequential:** This setting trades in the day-ahead market without anticipating the intraday market. Afterwards, the commitments are adjusted on the intraday market. To reflect this with our policy functions, we add the constraint $x_{t,h,q}^{ID} = 0\ \forall h, q$ to $X_t^{DA,\pi}(\cdot)$ and the constraint $x_{t,h}^{DA} = 0\ \forall h$ to $X_t^{ID,\pi}(\cdot)$.
- **DA+ID, integrated:** This is the basic setting considered in the remainder of the paper.
- **DA+ID, integrated, no arbitrage:** Integrated trading in both auction markets can speculate on the price difference between the hourly day-ahead market products and the price of the corresponding quarter-hours on the intraday market. In this setting, we modify the intraday market price process defined in Section 3.1 to enforce that the expected market prices are equal. Therefore, no arbitrage based on the expected price difference is possible. For this setting, we shift the expected intraday market prices to the expected day-ahead market prices and use them in the contribution function of the policy model (34):

$$\bar{P}_{t,t+1,h,q}^{ID'} := \bar{P}_{t,t+1,h,q}^{ID} - \frac{1}{4}\sum_{q'=0}^{3}\left(\bar{P}_{t,t+1,h,q'}^{ID} - \bar{P}_{t,t+1,h}^{DA}\right) \forall h \tag{40}$$

All settings are optimized with the *BADP-w* approach with seven days in the price history and evaluated with real-world data. Again, we run the simulation (estimating price processes, solving trading problem, and evaluation) 20 times, to capture the stochastics (see, Figure 5). We report the mean of the summed profits over the four seasons and the 95% confidence bands.

Table 4 compares our basic setting (DA+ID, integrated) with four other settings: only using the day-ahead market (DA), only using the intraday market (ID), using both markets with a

Table 4: Comparison to restricted trading settings

Trading and optimization	Profit €	Relative profit [%]	95% confidence band
Only day-ahead (DA)	37,693.59	39.75	[39.03, 40.47]
Only intraday (ID)	92,465.45	97.51	[97.20, 97.83]
DA+ID, sequential	87,221.54	91.98	[91.74, 92.22]
DA+ID, integrated	94,826.64	100.00	[99.61, 100.39]
DA+ID, integrated, no arbitrage	94,712.84	99.88	[99.39, 100.37]

sequential optimization (DA+ID, sequential), and trading integrated but without exploiting the price differences (DA+ID, integrated, no arbitrage). If only one market is used, the intraday market is much more profitable for the storage owner than the day-ahead market. This is because the intraday market prices fluctuate substantially more than the day-ahead market prices, which can be exploited by the fast ramping energy storage. The immense difference between the day-ahead and intraday profits decreases in the ramping time (not shown here). In the sequential day-ahead and intraday setting, the storage owner first trades on the day-ahead market, as if only the day-ahead market exists, i.e. exactly as in setting (DA). After this, the storage owner adjusts the commitments on the intraday market. The integrated bidding can enhance the profit by 8 percentage points compared to the sequential consideration of the markets and 2.5 percentage points compared to the intraday market setting. In our study, the intraday market setting outperforms the sequential setting. This is because decisions on the day-ahead market cause a bad starting point for trading on the intraday market and, thus, profit is lower. This changes with increasing ramping times (not shown here). If the storage owner does not exploit the expected price difference between the markets, the relative profit only decreases by 0.12 percentage points. The difference between these two settings is not significant. Thus, (almost) no arbitrage is possible, which indicates that the German auction markets are efficient.

7. Discussion

In this section, we discuss possible limitations of our research. We observed that integrated trading in both German auction markets outperforms sequential trading and trading in a single auction market. We did not investigate whether integrated trading in the auction markets outperforms trading in other markets (e.g., derivative markets, continuous intraday market, balancing markets, over-the-counter). Furthermore, we did not investigate if integrated trading with additional markets is beneficial. In the following, we discuss how this limits our research and how practice can still benefit.

Derivative markets: Trading derivatives can be an important instrument to hedge the risk (e.g., variance of returns), see Hain et al. (2018). Since we optimize expected profit, the derivative markets are not beneficial.

Continuous intraday market: As the continuous intraday market works similar to stock markets, a product's current price includes all relevant information and is therefore the best predictor for the future price of the product (see, e.g., Narajewski & Ziel (2020a)). Since the continuous intraday market starts with the intraday auction market clearing prices and both markets have the same temporal resolution, integrated trading in the intraday auction market and the continuous intraday market should not be profitable. Moreover, the continuous intraday market has a rather high market entrance barrier and a 24/7 trading floor is needed.

Nevertheless, sequential trading in the continuous intraday market after integrated trading in the auction markets should outperform pure auction market trading. Since the trading in the auction markets restricts the trading in the continuous intraday market, pure continuous intraday market trading could be even better than a heuristic solution of integrated trading.

Balancing markets: The balancing markets can be an important opportunity for fast ramping energy storage systems. We do not include them in our research for the following reason. The rules of the German balancing markets changed multiple times in recent years (e.g., 2018, 2019, and 2020). Since these changes happen during the time span considered in the numerical study of Section 6, nearly no data for the current regime exists.

Over-the-counter markets: For energy storage systems, the over-the-counter (OTC) markets are hardly relevant, because there, most energy is contracted several days or months ahead. Most energy storage systems are fully charged and discharged within hours, so they benefit most from short-term price fluctuations. Additionally, there is no public information about OTC markets available.

In summary, we are confident that we modeled the two most important German markets, but acknowledge that the integration of additional markets, like balancing markets, is desirable.

Based on our correspondence with practitioners and in line with literature, we see that market participants often trade in more than two markets but in a sequential manner. As seen in Section 6.5 this can decrease profit, even compared to trading in fewer markets. Nevertheless, market participants that trade in a sequential manner can use our model and approach to combine two links in the chain to improve their trades. This is the first step towards a fully anticipating trading strategy that integrates all relevant markets in an optimization problem.

8. Conclusion and insights

In this paper, we modeled the energy trading problem for both German auction markets with a competitive price forecast as a high-dimensional dynamic program. We implemented a suitable backwards ADP solution algorithm and showed that the approach successfully overcomes the resulting curse of dimensionality. Furthermore, the implemented approach outperforms an intuitive and widely used benchmark using an expectation model in a receding horizon fashion as well as a state-of-the-art approach from literature. However, it is important to use a long enough price history in the optimization's forecast to outperform the expectation model. This is an important observation as the resulting explosion of the DP's state space often forces researchers to use shorter histories, whereas this is much less a problem for expectation models. Furthermore, one has to use an approach that is designed for handling high-dimensional price processes, albeit common methods from literature cannot.

Regarding the integrated trading, we have shown that integrated trading in both German auction markets is beneficial. Further, we have shown that trading in a sequential manner can be even worse than only trading in the intraday auction market. This shows that storage owners that sequentially trade have to rethink this and anticipate later markets. This may also be true for markets not included in our research. To reduce the computational burden of such an integrated optimization it is sufficient to use self-schedule bidding in the auction markets. If one needs to reduce the computational burden even more, an expectation model with a competitive price process is more appropriate than an ADP approach with a simple price process.

Future work could go in several directions. First, from a methodological perspective, infinite-horizon dynamic programs should be considered. Whereas our proposed dynamic program is time-independent (except the day of the week), future work can reformulate the problem as a steady state problem and solve the storage problem with an infinite horizon. This has its own challenges. For example, the set of sample paths can no longer simply be sampled with the price processes because this depends upon the initial price history. Instead, one can use all past price histories of the last few years. Second, another research avenue would be to investigate the influence of ramping time and other properties of the energy storage on expected profit. Especially if investment costs are taken into account one can, for example, investigate the benefit of a small buffer battery for a pumped hydro storage to enable instant ramping.

References

Bonnans, J. F., Cen, Z., & Christel, T. (2011). Energy contracts management by stochastic programming techniques. *Annals of Operations Research*, *200*, 199–222.

Boomsma, T. K., Juul, N., & Fleten, S.-E. (2014). Bidding in sequential electricity markets: The nordic case. *European Journal of Operational Research*, *238*, 797–809.

Brijs, T., Geth, F., Jonghe, C. D., & Belmans, R. (2019). Quantifying electricity storage arbitrage opportunities in short-term electricity markets in the CWE region. *Journal of Energy Storage*, *25*, 100899.

Broadie, M., Glasserman, P., & Ha, Z. (2000). Pricing american options by simulation using a stochastic mesh with optimized weights. In *Nonconvex Optimization and Its Applications* (pp. 26–44). Springer US.

Crespo-Vazquez, J. L., Carrillo, C., Diaz-Dorado, E., Martinez-Lorenzo, J. A., & Noor-E-Alam, M. (2018). A machine learning based stochastic optimization framework for a wind and storage power plant participating in energy pool market. *Applied Energy*, *232*, 341–357.

Ding, H., Hu, Z., & Song, Y. (2015). Rolling optimization of wind farm and energy storage system in electricity markets. *IEEE Transactions on Power Systems*, *30*, 2676–2684.

Durante, J., Nascimento, J., & Powell, W. B. (2020). Backward approximate dynamic programming with hidden semi-markov stochastic models in energy storage optimization. Working paper, Princeton University, Princeton.

EPEX SPOT SE (2020). Market data. https://www.epexspot.com/en/market-data. Accessed 21 July 2020.

Fathima, A. H., & Palanisamy, K. (2015). Optimization in microgrids with hybrid energy systems – a review. *Renewable and Sustainable Energy Reviews*, *45*, 431–446.

Finnah, B. (2021). Optimal bidding functions for renewable energies in sequential electricity markets. *OR Spectrum*, . doi:10.1007/s00291-021-00646-9.

Finnah, B., & Gönsch, J. (2021). Optimizing trading decisions of wind power plants with hybrid energy storage systems using backwards approximate dynamic programming. *International Journal of Production Economics*, *238*, 108–155.

Fleten, S.-E., & Kristoffersen, T. K. (2008). Short-term hydropower production planning by stochastic programming. *Computers & Operations Research*, *35*, 2656–2671.

Ghavidel, S., Ghadi, M. J., Azizivahed, A., Aghaei, J., Li, L., & Zhang, J. (2020). Risk-constrained bidding strategy for a joint operation of wind power and CAES aggregators. *IEEE Transactions on Sustainable Energy*, *11*, 457–466.

Glasserman, P. (2003). *Monte Carlo Methods in Financial Engineering*. Springer New York.

Gönsch, J., & Hassler, M. (2016). Sell or store? an ADP approach to marketing renewable energy. *OR Spectrum*, *38*, 633–660.

Hain, M., Schermeyer, H., Uhrig-Homburg, M., & Fichtner, W. (2018). Managing renewable energy production risk. *Journal of Banking & Finance*, *97*, 1–19.

Hannah, L. A., & Dunson, D. B. (2011). Approximate dynamic programming for storage problems. In *Proceedings of the 28th International Conference on International Conference on Machine Learning* ICML'11 (pp. 337–344). USA: Omnipress.

Heredia, F. J., Cuadrado, M. D., & Corchero, C. (2018). On optimal participation in the electricity markets of wind power plants with battery energy storage systems. *Computers & Operations Research*, *96*, 316–329.

Jędrzejewski, A., Marcjasz, G., & Weron, R. (2021). Importance of the long-term seasonal component in day-ahead electricity price forecasting revisited: Parameter-rich models estimated via the LASSO. *Energies*, *14*, 3249.

Jiang, D. R., Pham, T. V., Powell, W. B., Salas, D. F., & Scott, W. R. (2014). A comparison of approximate dynamic programming techniques on benchmark energy storage problems: Does anything work? In *2014 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL)*. IEEE.

Khatami, R., Oikonomou, K., & Parvania, M. (2020). Look-ahead optimal participation of compressed air energy storage in day-ahead and real-time markets. *IEEE Transactions on Sustainable Energy*, *11*, 682–692.

Kim, J. H., & Powell, W. B. (2011). Optimal energy commitments with storage and intermittent supply. *Operations Research*, *59*, 1347–1360.

Kuznetsova, E., Ruiz, C., Li, Y.-F., & Zio, E. (2015). Analysis of robust optimization for decentralized microgrid energy management under uncertainty. *International Journal of Electrical Power & Energy Systems*, *64*, 815–832.

Löhndorf, N., & Minner, S. (2010). Optimal day-ahead trading and storage of renewable energies—an approximate dynamic programming approach. *Energy Systems*, *1*, 61–77.

Löhndorf, N., & Shapiro, A. (2019). Modeling time-dependent randomness in stochastic dual dynamic programming. *European Journal of Operational Research*, *273*, 650–661.

Löhndorf, N., Wozabal, D., & Minner, S. (2013). Optimizing trading decisions for hydro storage systems using approximate dual dynamic programming. *Operations Research*, *61*, 810–823.

Marcjasz, G. (2020). Forecasting electricity prices using deep neural networks: a robust hyper-parameter selection scheme. *Energies*, *13*, 4605.

Marcjasz, G., Uniejewski, B., & Weron, R. (2019). On the importance of the long-term seasonal component in day-ahead electricity price forecasting with narx neural networks. *International Journal of Forecasting*, *35*, 1520–1532.

Moghaddam, I. G., Nick, M., Fallahi, F., Sanei, M., & Mortazavi, S. (2013). Risk-averse profit-based optimal operation strategy of a combined wind farm–cascade hydro system in an electricity market. *Renewable Energy*, *55*, 252–259.

Narajewski, M., & Ziel, F. (2020a). Econometric modelling and forecasting of intraday electricity prices. *Commodity Markets*, *19*, 100–107.

Narajewski, M., & Ziel, F. (2020b). Ensemble forecasting for intraday electricity prices: Simulating trajectories. *Applied Energy*, *279*, 115801.

Narajewski, M., & Ziel, F. (2021). Optimal bidding on hourly and quarter-hourly day-ahead electricity price auctions. URL: https://arxiv.org/pdf/2104.14204.pdf.

Olaszi, B. D., & Ladanyi, J. (2017). Comparison of different discharge strategies of grid-connected residential pv systems with energy storage in perspective of optimal battery energy storage system sizing. *Renewable and Sustainable Energy Reviews*, *75*, 710–718.

Powell, W. B. (2019). A unified framework for stochastic optimization. *European Journal of Operational Research*, *275*, 795–821.

Rahman, H. A., Majid, M. S., Jordehi, A. R., Kim, G. C., Hassan, M. Y., & Fadhl, S. O. (2015). Operation and control strategies of integrated distributed energy resources: A review. *Renewable and Sustainable Energy Reviews*, *51*, 1412–1420.

Salas, D. F., & Powell, W. B. (2018). Benchmarking a scalable approximate dynamic programming algorithm for stochastic control of grid-level energy storage. *INFORMS Journal on Computing*, *30*, 106–123.

Scott, P. (2019). Approximate dynamic programming for energy storage with new results on instrumental variables and projected bellman errors. *Submitted to Operations Research*, .

Steffen, B. (2012). Prospects for pumped-hydro storage in germany. *Energy Policy*, *45*, 420–429.

Sterner, M., & Stadler, I. (2017). *Energiespeicher - Bedarf, Technologien, Integration*. Springer Berlin Heidelberg.

Uniejewski, B., & Weron, R. (2018). Efficient forecasting of electricity spot prices with expert and lasso models. *Energies*, *11*, 2039.

Vahedipour-Dahraie, M., Rashidizadeh-Kermani, H., Anvari-Moghaddam, A., & Siano, P. (2020). Risk-averse probabilistic framework for scheduling of virtual power plants considering demand response and uncertainties. *International Journal of Electrical Power & Energy Systems*, *121*, 106126.

Viehmann, J. (2017). State of the german short-term power market. *Zeitschrift für Energiewirtschaft*, *41*, 87–103.

Weitzel, T., & Glock, C. H. (2018). Energy management for stationary electric energy storage systems: A systematic literature review. *European Journal of Operational Research*, *264*, 582–606.

Weron, R. (2014). Electricity price forecasting: A review of the state-of-the-art with a look into the future. *International Journal of Forecasting*, *30*, 1030–1081.

Wozabal, D., & Rameseder, G. (2020). Optimal bidding of a virtual power plant on the spanish day-ahead and intraday market for electricity. *European Journal of Operational Research*, *280*, 639–655.

Xu, X., Hu, W., Cao, D., Huang, Q., Liu, Z., Liu, W., Chen, Z., & Blaabjerg, F. (2020). Scheduling of wind-battery hybrid system in the electricity market using distributionally robust optimization. *Renewable Energy*, *156*, 47–56.

Zéphyr, L., & Anderson, C. L. (2017). Stochastic dynamic programming approach to managing power system uncertainty with distributed storage. *Computational Management Science*, *15*, 87–110.

Zhou, Y. H., Scheller-Wolf, A., Secomandi, N., & Smith, S. (2016). Electricity trading and negative prices: Storage vs. disposal. *Management Science*, *62*, 880–898.

Zhu, D., Wang, Y., Yue, S., Xie, Q., Pedram, M., & Chang, N. (2013). Maximizing return on investment of a grid-connected hybrid electrical energy storage system. In *2013 18th Asia and South Pacific Design Automation Conference (ASP-DAC)* (pp. 638–643).

Ziel, F. (2016). Forecasting electricity spot prices using lasso: On capturing the autoregressive intraday structure. *IEEE Transactions on Power Systems*, *31*, 4977–4987.

Appendix A. Self-schedule bidding

To check the self-schedule assumption Assumption 2, we show numerically that the gains from price-volume bids are negligible for electricity prices with highly correlated noise (like in Germany). For this, we solve a simplified energy trading problem numerically with price-volume bids and with self-schedule bids and evaluate these with common out-of-sample scenarios. In this simplified setting, the PHS only trades in the day-ahead market and ramping is ignored. The notation is summarized in Appendix Appendix D. To show the effect of the correlation of the noise, we draw the scenarios based on the price process defined in Section 3.1, but artificially reduce the correlation. We denote the artificial correlation matrix as $\Sigma^{DA,\gamma}$. This is defined as

$$\Sigma^{DA,\gamma} = (\sigma_{i,j}^{DA,\gamma})_{i=1,\cdots,24;j=1,\cdots,24} := \begin{cases} \sigma_{i,j}^{DA} & i = j \\ \gamma\sigma_{i,j}^{DA} & \text{else} \end{cases} \tag{A.1}$$

with $\gamma \in [0, 1]$ and the originally estimated correlation matrix $\Sigma^{DA} = (\sigma_{i,j}^{DA})_{i=1,\cdots,24;j=1,\cdots,24}$. For $\gamma = 1$, the matrix $\Sigma^{DA,\gamma}$ has the correlation estimated with real-world data. For $\gamma = 0$, the correlation is zero.

We model the simplified decision problem as a two-stage stochastic program. Based on a state $S_t^{DA} = \left(R_t^{start}, x_t^{start}, P_t^{history}\right)$, the here-and-now decisions are B selling and buying bids for each hour. The wait-and-see decisions are the actual dispatching of the PHS after the market is cleared. The volume components of the bids are denoted $x_{t,h,b}^{buy,bid}$ and $x_{t,b,h}^{sell,bid}$ with hour h and the number of the bid $b \in \{1, \cdots, B\}$. The price components are denoted $P_{t,h,b}^{buy,bid}$ and $P_{t,h,b}^{sell,bid}$. While we model the volumes as continuous decision variable, the prices are fixed parameters. Therefore, we map each price parameter $P_{t,h,b}^{buy,bid}$ and $P_{t,h,b}^{sell,bid}$ a to volume $x_{h,b}^{buy,bid}$ and $x_{b,h}^{sell,bid}$.

As we replace the stochastic program with the SAA based on S scenarios, we sample the price scenarios $P_{t+1,h,s}^{DA}$ for each hour h based on the price process defined in Section 3.1 with correlation matrix $\Sigma^{DA,\gamma}$ and the price history $P_t^{history}$. Based on these samples, we can pre-process the parameters $a_{t,h,s,b}^{buy}$ and $a_{t,h,s,b}^{sell}$.

$$a_{t,h,s,b}^{buy} = \begin{cases} 1 & , P_{t,h,b}^{buy,bid} \geq P_{t+1,h,s}^{DA} \\ 0 & , \text{otherwise} \end{cases} \tag{A.2}$$

$$a_{t,h,s,b}^{sell} = \begin{cases} 1 & , P_{t,h,b}^{sell,bid} < P_{t+1,h,s}^{DA} \\ 0 & , \text{otherwise} \end{cases} \tag{A.3}$$

If $a_{t,h,s,b}^{buy} = 1$, the buying bid b in hour h and scenario s is accepted (analogously for $a_{t,h,s,b}^{sell}$).

$$\max \frac{1}{S}\sum_{s=1}^{S}\Delta t^{DA}\sum_{h=1}^{24} -\left(P^{DA}_{t+1,h,s}+c^{gf}\right)x^{buy}_{t,h,s}+P^{DA}_{t+1,h,s}x^{sell}_{t,h,s}+Q^{over}_{t+1,h,s}x^{over}_{t,h,s}-Q^{under}_{t+1,h,s}x^{under}_{t,h,s} \tag{A.4}$$

$$-c^{pump}z^{pump}_{t,h,s}-c^{turbine}z^{turbine}_{t,h,s}$$

s.t.

$$x^{buy}_{t,h,s}=\sum_{b=1}^{B}a^{buy}_{t,h,s,b}x^{buy,bid}_{t,h,b} \tag{A.5}$$

$$x^{sell}_{t,h,s}=\sum_{b=1}^{B}a^{sell}_{t,h,s,b}x^{sell,bid}_{t,h,b} \tag{A.6}$$

$$x^{pump}_{t,h,s}+x^{over}_{t,h,s}=x^{buy}_{t,h,s} \tag{A.7}$$

$$x^{turbine}_{t,h,s}+x^{under}_{t,h,s}=x^{sell}_{t,h,s} \tag{A.8}$$

$$R_{t,h,s}=R_{t,h-1,s}+\eta^{pump}\Delta t^{DA}x^{pump}_{t,h,s}-\frac{\Delta t^{DA}}{\eta^{turbine}}x^{turbine}_{t,h,s} \tag{A.9}$$

$$x^{comp}_{min}y^{comp}_{t,h,s}\le x^{comp}_{t,h,s}\le x^{comp}_{max}y^{comp}_{t,h,s}\quad \forall comp\in\{pump,turbine\} \tag{A.10}$$

$$z^{comp}_{t,h,s}\ge y^{comp}_{t,h,s}-y^{comp}_{t,h-1,s}\quad \forall comp\in\{pump,turbine\} \tag{A.11}$$

$$R_{t,h,s}\in[0,R_{max}] \tag{A.12}$$

$$x^{buy}_{t,h,b},x^{sell}_{t,h,b}\ge 0 \tag{A.13}$$

$$x^{over}_{t,h,s},x^{under}_{t,h,s},z^{pump}_{t,h,s},z^{turbine}_{t,h,s}\ge 0 \tag{A.14}$$

$$y^{pump}_{t,h,s},y^{turbine}_{t,h,s}\in\{0,1\} \tag{A.15}$$

All constraints are defined for all $s \in \{1,\cdots S\}$ and $h \in \{0,\cdots,23\}$. Only constraint (A.13) is defined for all $b \in \{1,\cdots B\}$ and $h \in \{0,\cdots,23\}$. Additionally, $R_{t,-1,s}$ is defined as R^{start}_t and $y^{pump}_{t,0,-1}=1_{x^{start}_t>0}$ and $y^{turbine}_{t,0,-1}=1_{x^{start}_t<0}$. The objective (A.4) maximizes the average profit over all samples. While pumping, the PHS must pay the market price $P^{DA}_{t+1,h,s}$ and the grid fees c^{gf}. While generating energy, the PHS receives the market price. If the PHS does not deliver the energy sold (e.g., empty storage), it has to pay the balancing price for under supply $Q^{under}_{t+1,h,s}$. If the PHS stores less energy than bought (e.g., full storage), it receives the balancing price for over supply $Q^{over}_{t+1,h,s}$. The increased maintenance costs during start-up are reflected by c^{pump} and $c^{turbine}$. Constraints (A.5) and (A.6) derive the contracted commitments based on the decided volumes $x^{(\cdot,bid)}_{t,h,b}$ which are mapped to the parameters $P^{(\cdot,bid)}_{t,h,b}$. Therefore, in scenario s and hour h, the PHS buys $x^{buy}_{t,h,s}$ and sells $x^{sell}_{t,h,s}$. Constraints (A.7) and (A.8) derive the difference between the contracted energy and the stored/delivered energy. Constraint (A.9) is the storage balance equation. Constraint (A.10) limits the power of the turbine and the pump to the feasible range. The binary variables $y^{pump}_{t,h,s}$ and $y^{turbine}_{t,h,s}$ are equal to one, if the pump/turbine work in scenario s and hour h. Constraint (A.11) captures if the pump or the turbine has to start-up. The constraints (A.12) to (A.15) define the domain of the variables.

If not stated otherwise, we use the same parameters as in the numerical study of Section 6. Since this section focuses on daily energy trading, we use $R_t^{start} = 0$ and $x_t^{start} = 0$. Moreover, we define balancing prices for over- and undersupply that are linear in the day-ahead market price.

$$Q_{t+1,h,s}^{over} = Q_{intercept}^{over} + Q_{slope}^{over} P_{t+1,h,s}^{DA} \tag{A.16}$$

$$Q_{t+1,h,s}^{under} = Q_{intercept}^{under} + Q_{slope}^{under} P_{t+1,h,s}^{DA} \tag{A.17}$$

The following parameters differ for the energy trading problem with price-volume bids and the problem with self-schedule bids:

- **Price-volume bids:** Here we use $S = 50$ price scenarios drawn with the estimated price process but artificial correlation matrix $\Sigma^{DA,\gamma}$. As parameters for the price components of the $B = 33$ bids, we use $P_{t,h,b}^{buy,bid}, P_{t,h,b}^{sell,bid}, \in \left\{-\infty, \bar{P}_{t,t+1,h}^{DA} - 30, \bar{P}_{t,t+1,h}^{DA} - 28, \cdots, \bar{P}_{t,t+1,h}^{DA} + 30, \infty\right\} \forall h$. Thus, self-schedule bids within ±30 €/MWh of the expected market prices are used.
- **Self-schedule bids:** Here we use a single price scenario ($S = 1$) with $P_{t+1,h,s}^{DA} = \bar{P}_{t,t+1,h}^{DA} \forall h$. Therefore, the prices are equal to the expected day ahead market on day $t + 1$ only using data known at day t. To force (A.4) to (A.15) to use self schedule bids, we use as price components of the bids $P_{t,h,b}^{buy,bid} = \infty \forall h$ and $P_{t,h,b}^{sell,bid} = -\infty \forall h$. Thus, only one bid ($B = 1$) results.

Obviously we cannot report the objective value of (A.4) but have to evaluate this with out-of-sample random numbers. In the following pseudocode we state how to apply (A.4) to (A.15) on out-of-sample data:

Algorithm 3 Evaluate (A.4) to (A.15)

1: Sample S day-ahead market prices $P_{t+1,h,s}^{DA}$
2: Solve (A.4) to (A.15)
3: Use decisions $x_{t,h,b}^{buy,bid}$ and $x_{t,h,b}^{sell,bid}$ and observe $P_{t+1,h}^{DA}$
4: Solve (A.4) to (A.15) with the single scenario $P_{t+1,h,s}^{DA} = P_{t+1,h}^{DA}$ and fixed $x_{t,h,b}^{buy,bid}$ and $x_{t,h,b}^{sell,bid}$

To show the effect of the correlation of the noise, we solve the daily trading problem for each of the 120 days in the test data presented in Section 6 with $\gamma \in \{1, 0.75, 0.5, 0.25, 0\}$. Since we estimated the price processes 20 times with cross validation, we solve the problem 20 times for each day. In contrast to the remainder of the paper, we only use the corresponding price history $P_t^{history}$ and do not evaluate with real-world data. The reason for this is that the real world data follows a price process with $\gamma = 1$. Therefore, we have to evaluate with randomly drawn scenarios based on the correlation matrix $\Sigma^{DA,\gamma}$. This procedure introduces a new level of randomness (evaluation level) which is not included in Figure 5. To reduce the computational burden, we set a relative MILP solution gap of 1% while deriving the first-stage decisions.

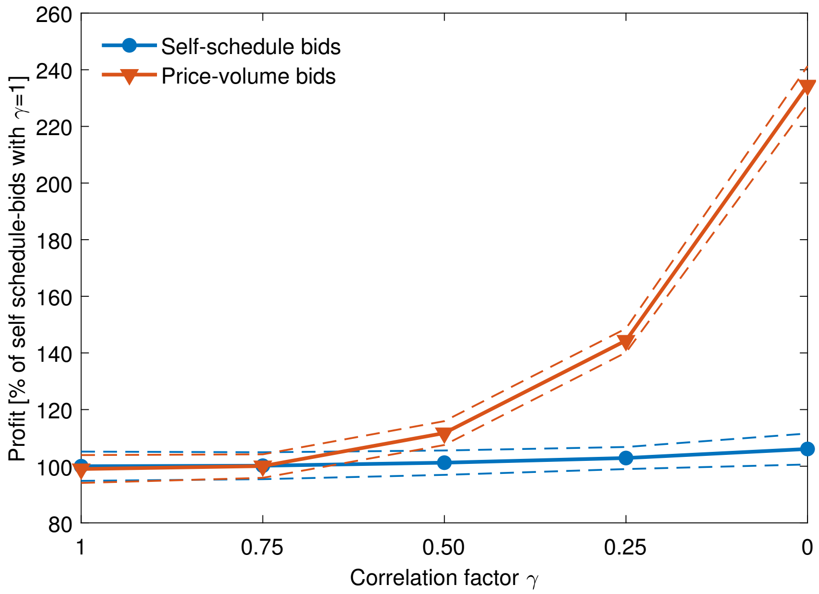

Figure A.9: Influence of correlation on expected profit ($\gamma = 0$: no correlation, $\gamma = 1$: correlation of German market)

Figure A.9 shows the relative average profit summed over the 120 days including the 95% confidence bands based on self-schedule bids and price-volume bids with $\gamma \in \{1, 0.75, 0.5, 0.25, 0\}$. For the correlation based on real-world data ($\gamma = 1$) there is no significant difference. The reason for this is, that energy storage systems are profitable if the price difference between the buying hours prices and selling hours prices is large enough. Since the forecast errors of the 24 day-ahead market prices are very highly correlated (coefficient of correlation up to 0.98 for consecutive hours), forecasts usually over- or underestimate the whole day's prices, which hardly influences the daily price differences. This justifies our self-schedule Assumption 2. Furthermore, the machine time with price-volume bids is over 5,000 times higher than the machine time with self-schedule bids (not shown here). With decreasing γ and therefore decreasing correlation the price-volume bids become increasingly profitable. For uncorrelated noise ($\gamma = 0$), price-volume bids dramatically outperform self-schedule bids by over 130 percentage points. Therefore, in energy markets with lower correlation, energy storage systems must use price-volume bids. For example, the Nordic markets behave very different than the German markets because of the immense penetration of storeable hydropower.

Please note that we only show above that price-volume bids are not needed for the daily energy trading. However, we are confident that this is sound enough for the following reasons: First, from a theoretical perspective, price-volume bids outperform self-schedule bids if one is able so solve the corresponding problems exactly. Thus, it is reassuring that price-volume bids never perform significantly worse than self-schedule bids. Second, one may ask whether this also holds for the problem with multiple days. We did not investigate this setting here, as the tractability mentioned above becomes more difficult as the problem grows. However, since most energy storage systems are fully charged/discharged within hours, the daily price differences are the most important ones for energy storage systems. In contrast, hydropower plants have huge reservoirs (discharge time of days to weeks) and a natural water inflow. For these power plants it is essential to release the exogenous resource at high price hours to be profitable (e.g. Fleten & Kristoffersen (2008), Boomsma et al. (2014)).

Appendix B. Expectation model versus stochastic program

To demonstrate that the policy designed in Section 4 works well, we compare it with a two-stage stochastic program in a single-day setting. While the here-and-now decisions correspond to the day-ahead market stage, the wait-and-see decisions are the intraday market stage decisions. To solve the two-stage stochastic program we use the sample average approximation with S scenarios.

$$\begin{aligned}
&\max \frac{1}{S}\sum_{s=1}^{S}\sum_{h=0}^{23} -\Delta t^{DA} P^{DA}_{t+1,h,s} x^{DA}_{t,h} + \sum_{q=0}^{3} -\Delta t^{ID} \bar{P}^{ID}_{t+1,h,q,s} x^{ID}_{t,h,q,s} - \Delta t^{ID} c^{gf} x^{pump}_{t,h,q,s} \\
&+ \frac{\Delta t^{ramp}}{2}\left(\left(\Delta x^{pump,up}_{t,h,q,s} + \Delta x^{turbine,down}_{t,h,q,s}\right)\bar{Q}^{over}_{t+1,h,q,s} - \left(\Delta x^{pump,down}_{t,h,q,s} + \Delta x^{turbine,up}_{t,h,q,s}\right)\bar{Q}^{under}_{t+1,h,q,s}\right) \\
&- c^{pump} z^{pump}_{t,h,q,s} - c^{turbine} z^{turbine}_{t,h,q,s}
\end{aligned} \tag{B.1}$$

s.t.

$$x^{DA}_{t,h} \in \left[-x^{turbine}_{max}, x^{pump}_{max}\right] \tag{B.2}$$

$$x^{DA}_{t,h} + x^{ID}_{t,h,q,s} = x^{pump}_{t,h,q,s} - x^{turbine}_{t,h,q,s} \tag{B.3}$$

$$x^{comp}_{min} y^{comp}_{t,h,q,s} \leq x^{comp}_{t,h,q,s} \leq x^{comp}_{max} y^{comp}_{t,h,q,s} \quad \forall comp \in \{pump, turbine\} \tag{B.4}$$

$$y^{pump}_{t,h,q,s} + y^{turbine}_{t,h,q,s} \leq 1 \tag{B.5}$$

$$\Delta x^{comp,up}_{t,h,q,s} - \Delta x^{comp,down}_{t,h,q,s} = x^{comp}_{t,h,q,s} - x^{comp}_{t,h,q-1,s} \quad \forall comp \in \{pump, turbine\} \tag{B.6}$$

$$\begin{aligned}
R_{t,h,q,s} = R_{t,h,q-1,s} &+ \eta^{pump}\left(\Delta t^{ID} x^{pump}_{t,h,q,s} - \frac{\Delta t^{ramp}}{2}\Delta x^{pump,up}_{t,h,q,s} + \frac{\Delta t^{ramp}}{2}\Delta x^{pump,down}_{t,h,q,s}\right) \\
&- \frac{1}{\eta^{turbine}}\left(\Delta t^{ID} x^{turbine}_{t,h,q,s} - \frac{\Delta t^{ramp}}{2}\Delta x^{turbine,up}_{t,h,q,s} + \frac{\Delta t^{ramp}}{2}\Delta x^{turbine,down}_{t,h,q,s}\right)
\end{aligned} \tag{B.7}$$

$$R_{t,h,q,s} \in [0, R_{max}] \tag{B.8}$$

$$y^{pump}_{t,h,q,s}, y^{turbine}_{t,h,q,s} \in \{0,1\} \tag{B.9}$$

$$x^{pump}_{t,h,q,s}, x^{turbine}_{t,h,q,s}, \Delta x^{pump,up}_{t,h,q,s}, \Delta x^{pump,down}_{t,h,q,s}, \Delta x^{turbine,up}_{t,h,q,s}, \Delta x^{turbine,down}_{t,h,q,s} \geq 0 \tag{B.10}$$

$$y^{comp}_{t,h,q,s} - y^{comp}_{t,h,q-1,s} \le z^{comp}_{t,h,q,s} \quad \forall comp \in \{pump, turbine\} \tag{B.11}$$
$$z^{pump}_{t,h,q,s}, z^{turbine}_{t,h,q,s} \ge 0 \tag{B.12}$$

All constraints are defined for all $q \in \{0, 1, 2, 3\}$, $h \in \{0, \cdots, 23\}$, and $s \in \{1, \cdots, S\}$. Additionally, $(\cdot)_{t,h,4,s}$ is defined as $(\cdot)_{t,h+1,0,s}$ and $R_{t,0,-1,s} = R^{start}_t$ and $x^{pump}_{t,0,-1,s} = \max\{x^{start}_t, 0\}$ and $x^{turbine}_{t,0,-1,s} = \max\{-x^{start}_t, 0\}$. Since this section focuses on daily energy trading, we use $R^{start}_t = 0$ and $x^{start}_t = 0$.

The SAA (B.1) to (B.12) is the stochastic counterpart of the policy model. The scenarios for the day-ahead market price $P^{DA}_{t+1,h,s}$ are sampled based on the price process defined in Section 3.1 and the current price history. The samples for the expected intraday market price $\bar{P}^{ID}_{t+1,h,q,s}$ are derived based on the price process defined in Section 3.1 conditioned on the sample $P^{DA}_{t+1,h,s}$. The expected balancing prices are computed based on these the usual way. Obviously, we cannot report the objective value of (B.1) but have to evaluate this with the real-world data in our test data set. In the following pseudocode we state how to apply (B.1) to (B.12) to real-world data:

Algorithm 4 Evaluate (B.1) to (B.12)

1: sample S day-ahead market prices $P^{DA}_{t+1,h,s}$ and compute the conditioned expectations $\bar{P}^{ID}_{t+1,h,q,s}$, $\bar{Q}^{over}_{t+1,h,q,s}$, and $\bar{Q}^{under}_{t+1,h,q,s}$
2: solve (B.1) to (B.12)
3: use decisions $x^{DA}_{t,h}$ and observe $P^{DA}_{t+1,h}$
4: compute the expectations $\bar{P}^{ID}_{t+1,h,q}$, $\bar{Q}^{over}_{t+1,h,q}$, and $\bar{Q}^{under}_{t+1,h,q}$ conditioned on the observation $P^{DA}_{t+1,h}$
5: solve (B.1) to (B.12) with the single scenario $\bar{P}^{ID}_{t+1,h,q}$, $\bar{Q}^{over}_{t+1,h,q}$, and $\bar{Q}^{under}_{t+1,h,q}$ and fixed $x^{DA}_{t,s}$.
6: use decisions $x^{DA}_{t,h,q}$ and observe $P^{ID}_{t+1,h,q}$

We solve the daily problem for each day in our test data set and evaluate the decisions with the real world data. We generate the samples based on the price processes with seven days in the price history. Since we estimated the price processes 20 times, we solve the daily problems 20 times and report the relative mean profits for $S \in \{10, 20, 30, 40, 50\}$ and the 95% confidence bands. To reduce the computational burden, we set a relative MILP solution gap of 1% while deriving the first-stage decisions (Step 2 in Algorithm 4).

Figure B.10 shows the relative cumulative (over all days) profit and the relative cumulative machine time of the stochastic program with different S. The relative profit slightly increases in the number of scenarios, but even with 50 scenarios, the stochastic program cannot significantly outperform the substantially faster and simpler expectation model. In addition, the expectation model has a much smaller confidence band. The width of the confidence band of the stochastic program decreases in S. Since the machine time dramatically increases in S, we do not test higher values. Even with $S = 50$, the stochastic program is 400 times slower than the expectation model. This shows that the daily expectation model works well. It would be possible to learn value

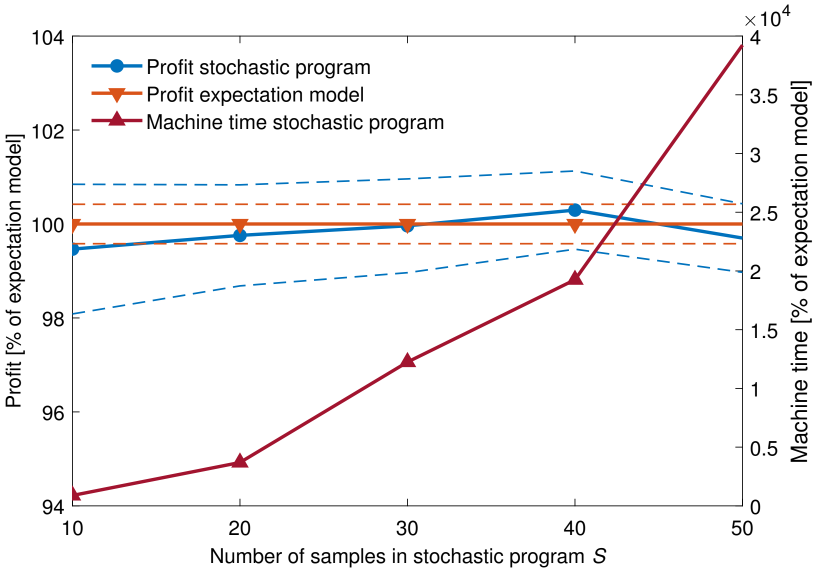

Figure B.10: Influence of the number of scenarios on the expected profit

functions with *BADP-w* and the policies (expectation model) designed in Section 4 and derive the decisions while evaluating with the stochastic program, but we do not think that the stochastic program would benefit dramatically enough to significantly outperform the expectation model.

Appendix C. Evaluating BADP

Using the value function approximation derived by *BADP* to evaluate the induced policies works similar to Step 4, Step 11, and Step 14 in Algorithm 1. The execution logic can be found in Algorithm 5.

Algorithm 5 Evaluate BADP

1: initialize S_0^{DA}
2: **for** $t = 0$ to $T-1$ **do**
3: compute expected price history $\bar{P}_{t,t+1}^{history} := \mathbb{E}\left[P_{t+1}^{history} | P_t^{history}\right]$
4: $(\bar{\rho}_1, \cdots, \bar{\rho}_N) := \arg\min\limits_{\rho_1, \cdots, \rho_N} \left\| \bar{P}_{t,t+1,n}^{history} - \sum\limits_{k=1}^{N} \rho_k P_{t+1,k}^{history} \right\|_2$
5: derive $x_t^{DA} := X_t^{DA,\pi}(S_t^{DA})$ based on $(\bar{\rho}_1, \cdots, \bar{\rho}_N)$
6: use decisions x_t^{DA} and observe P_{t+1}^{DA}
7: compute expected price history $\bar{P}_{t,t+1}^{history} := \mathbb{E}\left[P_{t+1}^{history} | P_t^{history}, P_{t+1}^{DA}\right]$
8: $(\bar{\rho}_1, \cdots, \bar{\rho}_N) := \arg\min\limits_{\rho_1, \cdots, \rho_N} \left\| \bar{P}_{t,t+1,n}^{history} - \sum\limits_{k=1}^{N} \rho_k P_{t+1,k}^{history} \right\|_2$
9: derive $x_t^{ID} := X_t^{ID,\pi}(S_t^{DA})$ based on $(\bar{\rho}_1, \cdots, \bar{\rho}_N)$
10: use decisions x_t^{ID} and observe P_{t+1}^{ID}
11: **end for**

Algorithm 5 shows how we use the value function approximation derived with *BADP* to derive decisions for the day-ahead and the intraday market. More precisely, how we evaluate the policy functions $X_t^{DA,\pi}\left(S_t^{DA}\right)$ and $X_t^{ID,\pi}\left(S_t^{ID}\right)$. Step 1 defines the initial state S_0^{DA}. Step 2 iterates over the optimization horizon. In Step 3, we derive the next day's expected price history conditioned on the current price history. The expected price history is regressed on the N sampled price histories (from Algorithm 1) to derive the regression parameters $(\bar{\rho}_1, \cdots, \bar{\rho}_N)$ in Step 4. These are used to compute $X_t^{DA,\pi}\left(S_t^{DA}\right) = \arg\max\limits_{x_t^{DA}} \max\limits_{x_t^{ID}} \left\{ C_t^{policy}\left(S_t^{DA}, \left(x_t^{DA}, x_t^{ID}\right)\right) + \sum\limits_{k=1}^{N} \bar{\rho}_k \bar{V}_{t+1}^{policy}\left(S_{t+1}^{DA} | P_{t+1,k}^{history}\right)\right\}$ s.t. (9), (11) to (18), (23) to (24), (33). Afterwards, the decision x_t^{DA} is applied and the MCPs P_{t+1}^{DA} becomes known. With the realized day-ahead market prices, the expected price history $\bar{P}_{t,t+1}^{history}$ is updated (Step 7) and again regressed on the sampled price histories (Step 8). The decision for the intraday market is $X_t^{ID,\pi}\left(S_t^{ID}\right) = \arg\max\limits_{x_t^{ID}} \mathbb{E}\left[C_t^{ID}\left(S_t^{ID}, x_t^{ID}, W_{t+1}^{ID}\right) \middle| S_t^{ID}\right] + + \sum\limits_{k=1}^{N} \bar{\rho}_k \bar{V}_{t+1}^{policy}\left(S_{t+1}^{DA} | P_{t+1,k}^{history}\right)$ s.t. (11) to (18), (23) to (24), (26). The decision x_t^{ID} is used and the intraday market prices become known.

Appendix D. Notation and parameter values

Table D.5: Notation introduced in Section 3

Indices	
t	Day
h	Hour
q	Quarter-hour
Endogenous variables	
$x_t^{DA} = \left(x_{t,0}^{DA}, \cdots, x_{t,23}^{DA}\right)$	Day-ahead market decision
$x_t^{ID} = \left(x_{t,0,0}^{ID}, \cdots, x_{t,23,3}^{ID}\right)$	Intraday market decision
S_t^{DA}	State variable in the day-ahead market stage
S_t^{ID}	State variable in the intraday market stage
R_t^{start}	Storage level at the start of day t
x_t^{start}	Inflow/outflow volume at the start of day t
$x_{t,h,q}^{turbine}$,$x_{t,h,q}^{pump}$	Inflow/outflow volume
$y_{t,h,q}^{turbine}$,$y_{t,h,q}^{pump}$	Binary working mode variable
$z_{t,h,q}^{turbine}$,$z_{t,h,q}^{pump}$	Start-up variable
$\Delta x_{t,h,q}^{turbine,up}$,$\Delta x_{t,h,q}^{turbine,down}$	Up/down ramping volume of the turbine
$\Delta x_{t,h,q}^{pump,up}$,$\Delta x_{t,h,q}^{pump,down}$	Up/down ramping volume of the pump
$R_{t,h,q}$	Storage level
Exogenous information	
$P_t^{DA} = \left(P_{t,1}^{DA}, \cdots, P_{t,24}^{DA}\right)$	Day-ahead market prices
$P_t^{ID} = \left(P_{t,0,0}^{ID}, \cdots, P_{t,23,3}^{ID}\right)$	Intraday market prices
$P_t^{history}$	Past prices used for price forecast
$Q_t^{over} = \left(Q_{t,0,0}^{over}, \cdots, Q_{t,23,3}^{over}\right)$	Payments for overproduction
$Q_t^{under} = \left(Q_{t,0,0}^{under}, \cdots, Q_{t,23,3}^{under}\right)$	Payments for underproduction
W_t^{DA}	Exogenous information which become known after the day ahead market stage
W_t^{ID}	Exogenous information which become known after the intraday market stage

Parameters	
T	Number of time periods
$DoW_{(\cdot)}^{(\cdot)}$	Day of the week dummy
$\beta_{(\cdot)}^{(\cdot)}$	Regression parameters in price forecast
$\Sigma^{(\cdot)}$	Correlation matrix
$\epsilon_t^{(\cdot)}$	Noise
x_{min}^{pump},x_{max}^{pump}	Working range of the pump
$x_{min}^{turbine}$,$x_{max}^{turbine}$	Working range of the turbine
R_{max}	Storage capacity
Δt^{DA}	Temporal resolution of the day-ahead market
Δt^{ID}	Temporal resolution of the intraday market
Δt^{ramp}	Ramping time
η^{pump},$\eta^{turbine}$	Efficiency of the pump/turbine
c^{gf}	Grid fees
c^{pump},c^{pump}	Start-up cost of the pump/turbine
$Q_{intercept}^{over}, Q_{intercept}^{under}$	Intercept of the linear imbalance cost
$Q_{slope}^{over}, Q_{slope}^{under}$	Slope of the linear imbalance cost
Functions	
C_t^{DA}	One-stage profit of the day-ahead market stage
C_t^{ID}	One-stage profit of the intraday market stage
$S^{M,DA}$	Transition function from the day-ahead market stage to the intraday market stage
$S^{M,ID}$	Transition function from the intraday market stage to the day-ahead market stage
$X_t^{DA,\pi}$	Policy function of the day-ahead market stage
$X_t^{ID,\pi}$	Policy function of the intraday market stage
V_t^{DA}	Value function of the day-ahead market stage
V_t^{ID}	Value function of the intraday market stage

Table D.6: Notation introduced in Section 4

Endogenous variables	
x_t^{policy}	Decision variable of the policy model
Exogenous information	
W_t^{policy}	Exogenous information of the policy model
Parameters	
$\bar{P}^{DA}_{t,t+1,h}$	Expectation of the day-ahead market prices
$\bar{P}^{ID}_{t,t+1,h,q}$	Expectation of the intraday market prices
$\bar{Q}^{over}_{t,t+1,h,q}$	Expectation of the balancing prices for overproduction
$\bar{Q}^{under}_{t,t+1,h,q}$	Expectation of the balancing prices for underproduction
Functions	
C_t^{policy}	One-stage profit of the policy model
$S^{M,policy}$	Transition function of the policy model
V_t^{policy}	Value function of the policy model
$\bar{V}^{policy}$	Value function approximation of the policy model
P	Transition probabilities

Table D.7: Notation introduced in Section 5

Indices	
n	Sample index
Parameters	
$P^{history}_{t,n} = \left(P^{history,1}_{t,n}, \cdots, P^{history,840}_{t,n}\right)$	Sample of the price history
$\bar{P}^{history}_{t,t+1,n}$	Expectation of the next price history based on the n-th sample
$\bar{P}^{history}_{t,t+1}$	Expectation of the next price history based on the realized price history
$\bar{\rho}_n$	Interpolation/extrapolation parameters
$K^d_{t,k}$	Upper/lower bounds for the Interpolation/extrapolation parameters
w_t	Vector of weights
e_j	j-th Unit vector
d	Bandwidth
N	Number of sample paths

Table D.8: Notation introduced in Section 6

Parameters	
K	Number of iterations of *ADDP*
D	Number of days in the price history

Table D.9: Parameter values for the numerical study

Parameter	Value	Parameter	Value
R_{max}	100 MWh	T	30
x^{pump}_{min}	5 MW	$x^{turbine}_{min}$	5 MW
x^{pump}_{max}	10 MW	$x^{turbine}_{max}$	10 MW
Δt^{DA}	1 h	Δt^{ID}	15 min
Δt^{ramp}	2 min	c^{gf}	5 €/MWh
c^{pump}	15 €	$c^{turbine}$	15 €
η^{pump}	0.9	$\eta^{turbine}$	0.9
$Q^{over}_{intercept}$	−3 €/MWh	Q^{over}_{slope}	$\frac{1}{1.2}$
$Q^{under}_{intercept}$	3 €/MWh	Q^{under}_{slope}	1.2
N	50	c^{gf}	5 €/MWh

Table D.10: Notation introduced in Appendix Appendix A

Indices	
b	Bid
s	Scenario
Parameters	
B	Number of bids
S	Number of scenarios
γ	Correlation factor
$\Sigma^{DA,\gamma}$	Artificial correlation matrix
$P^{buy,bid}_{t,h,b}$	Price component of the buying bid
$P^{sell,bid}_{t,h,b}$	Price component of the selling bid
$a^{buy}_{t,h,s,b}$	Indicator if buying bid is accepted
$a^{sell}_{t,h,s,b}$	Indicator if selling bid is accepted
Endogenous variables	
$x^{buy,bid}_{t,h,b}$	Volume component of the buying bid
$x^{sell,bid}_{t,h,b}$	Volume component of the selling bid
$x^{buy}_{t,h,s}$	Contracted buying volume
$x^{sell}_{t,h,s}$	Contracted selling volume
$x^{over}_{t,h,s}$	Overproduction
$x^{under}_{t,h,s}$	Underproduction