
Quantile-SMPC for Grid-Interactive Buildings with Multivariate Temporal Fusion Transformers

Spencer Hutchinson*

University of California, Santa Barbara
Santa Barbara, CA
shutchinson@ucsb.edu

Abraham P. Vinod François G. Germain Stefano Di Cairano

Christopher R. Laughman Ankush Chakrabarty

Mitsubishi Electric Research Laboratories (MERL)

Cambridge, MA

{vinod,germain,dicairano,laughman,chakrabarty}@merl.com

Abstract

Exogenous disturbances such as occupancy and weather strongly affect the performance of grid-interactive buildings, making accurate probabilistic forecasting essential for robust control. We propose a new MQF2-TFT architecture that captures both temporal and cross-variable dependencies for multi-horizon disturbance forecasting. To integrate forecasts into stochastic MPC, we develop a quantile-based sample method that enforces chance constraints despite non-Gaussianity. Experiments with real office building data demonstrates improved key performance indicators over baselines.

1 Motivation

Exogenous disturbances such as occupancy and outdoor conditions have a significant effect on building thermal dynamics and grid interactivity, affecting key performance indicators (KPIs) such as energy consumption. To design performant building energy management systems (BEMS) based on predictive control paradigms, it is critical to obtain accurate forecasts of these disturbances. Probabilistic forecasting enables robust decision-making (as opposed to deterministic) by computing policies that are uncertainty-aware: aware of statistics around the mean/median disturbance forecast [25, 13, 28].

While a variety of probabilistic forecasting models have been proposed, ranging from Gaussian processes to deep autoregressive models, the temporal fusion transformer (TFT) has emerged as particularly well-suited for applications requiring medium-to-long horizon forecasts [10]. This is because the TFT combines attention-based mechanisms for capturing long-range temporal dependencies with gating mechanisms that enable selective conditioning on relevant historical inputs. These features are especially valuable for BEMS, where disturbance trajectories such as occupancy or solar irradiance exhibit strong temporal patterns at daily and weekly scales. However, the classical TFT formulation produces forecasts independently across disturbance variables, and therefore *cannot fully capture the strong inter-variable dependencies inherent to building dynamics*: for example, correlations between human occupants and internal heat gains, or between ambient temperature and cooling capacity. To address this limitation, we augment the TFT with the multivariate quantile function forecaster (MQF2) [7]. MQF2 enables generative multi-variate forecasts by learning the

*Research completed during an internship at MERL.

joint quantile function across all disturbance variables, thereby complementing the TFT. The TFT learns temporal dependencies, while MQF2 explicitly learns cross-variable dependencies at each horizon.

A major challenge arises when integrating these forecasts into the SMPC optimization problem. SMPC often relies on chance-constrained optimization to ensure that operational constraints are satisfied with high probability. Standard formulations of SMPC usually assume Gaussian disturbances or employ scenario-based sampling to ensure the underlying robust optimal control problem is tractable. However, oversimplifying assumptions can fail to capture the skewed and multi-modal distributions of building disturbances; furthermore, scenario-based methods exhibit prohibitive sampling requirements at week-long or month-long prediction horizons. Instead, quantile-based formulations organically allow explicit control over risk levels in chance constraints without requiring Gaussian assumptions [27]. Therefore, we adopt a quantile-based SMPC approach that is informed by the MQF2-TFT, which enforces chance constraints directly using scenarios auto-regressively generated from the forecast distribution. This preserves tractability while faithfully leveraging the expressive non-Gaussian forecasts. Further relevant works are listed in Appendix A.

Contributions. The main contribution of this work is to bridge transformer-based probabilistic forecasting with SMPC, enabling expressive disturbance models to directly inform GEB control. In addition, we propose: (i) a quantile-driven formulation using the MQF2-augmented TFT that captures complex, non-Gaussian disturbance distributions while avoiding restrictive Gaussian assumptions or large-scale scenario sampling; (ii) an autoregressive feedback mechanism within MQF2-TFT to generate consistent multivariate scenarios across long horizons; and, (iii) we validate the proposed SMPC framework on real-world grid-interactive efficient building (GEB) data, demonstrating reduced energy costs and occupant discomfort relative to Gaussian and vector autoregressive baselines.

2 Methodology

GEB Dynamics. We consider the problem of the control of a GEB with photovoltaics and battery storage; see Appendix B for more details. The control inputs $u_t = [P_{hp}, P_{bat_{chg}}, P_{bat_{chg}}]$ are the heat pump capacity and the charge/discharge power of the battery, respectively. Exogenous disturbances are denoted $w_t = [T_{amb}, I, P_{int}]$, which are the ambient temperature T_{amb} , solar irradiance I and internal (sensible) heat loads P_{int} , respectively. The control inputs and disturbances impact the internal temperature of the building T_{in} and the state-of-charge of the battery Q_{bat} . The grid is treated as an infinite bus, and interactivity is modeled by a conservation equation. The primary objectives of the controller are to keep the internal temperature T_{in} within a comfortable range $T_{in}^{min} \leq T_{in} \leq T_{in}^{max}$, while reducing the amount of energy bought from the grid: by relying on the GEB’s internal solar energy capture and storage mechanisms.

Forecasting Model. The goal of the forecasting model is to predict the distribution of the future disturbances conditioned on past disturbances and on known covariates. That is, it approximates the true conditional distribution,

$$\mathbb{P}(w_{t+1:t+H} | w_{1:t}, c_{1:t+H}) \approx f(w_{t+1:t+H}; w_{t-k:t}, c_{t-k:t+H}) \quad (1)$$

where $w_{t+1:t+H}$ are the future disturbances, $w_{t-k:t}$ are the past disturbances, $c_{t-k:t+H}$ are known covariates, H is the prediction window, and k is the look-back window. Note that $w_t \in \mathbb{R}_{\geq 0}^3$. Our approximator (1) combines the TFT architecture [10] with the MQF2 architecture [7]. As opposed to the TFT architecture alone, which only provides univariate quantile forecasts, this proposed architecture allows for generative multi-variate forecasting while maintaining the advantages of TFT. More precisely, we use the TFT architecture augmented with an *autoregressive feedback*, to provide a horizon-specific latent \mathbf{h}_τ , i.e.,

$$\mathbf{h}_\tau = F_\theta(w_{t-k:t}, c_{t-k:t+\tau}, \hat{w}_{t+1:t+\tau-1}, \tau), \quad \tau = 1, \dots, H,$$

where $\hat{w}_{t+1:t+\tau-1}$ are the forecasts from the previous horizons. A shared MQF2 conditional decoder then uses this latent \mathbf{h}_τ to generate a joint forecast sample across all disturbance variables, rather than treating each channel as independent. That is,

$$\hat{w}_{t+\tau} = g_\phi(\mathbf{h}_\tau, \alpha), \quad \alpha \sim U([0, 1]^m).$$

By combining autoregressive feedback with the MQF2 decoder, the model learns correlations both across variables and over time, enabling coherent multivariate trajectories. We refer to Appendix C

for additional details on the model and training hyperparameters, and to our comparative study demonstrating that AR-TFT-MQF2 generally exhibits better distributional learning than vanilla TFT.

Quantile-SMPC. We propose a SMPC policy for GEB control that uses the samples from the forecasting model to enforce constraint satisfaction. Standard SMPC approaches (e.g. [12]) choose the next control input by solving a chance-constrained program enforcing high-probability constraint satisfaction on future states, i.e.,

$$\min_{\hat{\mathbf{x}}, \mathbf{u}} \mathbb{E} [J(\hat{\mathbf{x}}, \mathbf{u})] \quad (2a)$$

$$\text{subject to : } \mathbb{P}(C\hat{x}_\tau \leq T_{\text{in}}^{\max}) \geq p, \quad \forall \tau \in [t, t+N] \quad (2b)$$

$$\mathbb{P}(C\hat{x}_\tau \geq T_{\text{in}}^{\min}) \geq p, \quad \forall \tau \in [t, t+N] \quad (2c)$$

$$\hat{x}_{\tau+1} = A\hat{x}_\tau + Bu_\tau + G\hat{w}_\tau, \quad \forall \tau \in [t, t+N-1], \hat{x}_\tau = x_t \quad (2d)$$

$$(\hat{w}_t, \dots, \hat{w}_{t+N-1}) \sim f, \quad (2e)$$

where x_t is the current state, \hat{x}_τ are the predicted states under the linear dynamics in (2d) and the probabilistic disturbance forecast $(\hat{w}_t, \dots, \hat{w}_{t+N-1}) \sim f$, J is the performance-based cost function, and the bold variables denote the concatenation at all time steps; $\hat{\mathbf{x}} = [\hat{x}_t^\top, \dots, \hat{x}_{t+N}^\top]$, $\mathbf{u} = [u_t^\top, \dots, u_{t+N}^\top]$. Given the linear dynamics, the internal temperature is encoded as $T_{\text{in}} = Cx_t$, and therefore (2b) enforces the temperature comfort constraints with high probability.

Unfortunately, we cannot solve (2) directly, because our forecasting model f does not have a tractable analytic representation and therefore we cannot compute the chance constraints (2b). To circumvent this, we use an approach that enforces the chance-constraints using only scenarios or samples drawn from the forecasting model. To illustrate our approach, we first point out that (2b) can be written explicitly in terms of the control inputs \mathbf{u} (see Appendix D.1 for details):

$$\mathbb{P}_{\hat{\mathbf{w}} \sim f}(S_\tau^x x_t + S_\tau^u \mathbf{u} + S_\tau^u \hat{\mathbf{w}} \leq T_{\text{in}}^{\max}) \geq p \text{ and } \mathbb{P}_{\hat{\mathbf{w}} \sim f}(S_\tau^x x_t + S_\tau^u \mathbf{u} + S_\tau^u \hat{\mathbf{w}} \geq T_{\text{in}}^{\min}) \geq p. \quad (3)$$

Then, leveraging the approach from [27], we draw samples from the forecasting model and use these to compute the empirical quantile function of $S_\tau^u \hat{\mathbf{w}}$, denoted by $\hat{Q}_{S_\tau^u \hat{\mathbf{w}}}(\cdot)$. This empirical quantile function is then used to define deterministic constraints that enforces (2b) conservatively,

$$S_\tau^x x_t + S_\tau^u \mathbf{u} + \hat{Q}_{S_\tau^u \hat{\mathbf{w}}}(\Delta) \leq T_{\text{in}}^{\max} \text{ and } S_\tau^x x_t + S_\tau^u \mathbf{u} + \hat{Q}_{S_\tau^u \hat{\mathbf{w}}}(1 - \Delta) \geq T_{\text{in}}^{\min}, \quad (4)$$

where Δ is an appropriately chosen probability threshold. In particular, we choose $\Delta = p + \sqrt{\frac{\ln(1/\beta)}{2M}}$ for M samples, which was shown by [27] to ensure that the empirical quantile constraint on \mathbf{u} in (4) is a sufficient condition for the exact chance constraints in (2b) with probability at least $1 - \beta$. The complete pseudocode for our proposed SMPC algorithm is provided in Appendix D.2.

3 Results

Baselines. We compare our proposed approach with the following SMPC baselines. The implementation details of these baselines are given in Appendix E. **Gaussian:** Computes the SMPC policy with the disturbances taken to be iid Gaussian, with empirical mean and covariance from historical data. **VAR:** Computes the SMPC policy with the disturbances taken to be generated by a vector autoregressive (VAR) model with the same context length as our proposed model. **Cantelli:** Uses the same forecasting model as the proposed approach, but enforces the chance constraints with Cantelli's inequality using the empirical mean and variance of the forecasting samples; c.f. [5, 16].

Experimental Setup. We use a linear thermal dynamics model for the GEB; this is detailed in Appendix B. A real disturbance dataset for 2021-2024 measured at the Mitsubishi Electric SUSTIE building is used for both training and as ground-truth disturbances for simulation. The SUSTIE building is a next-generation commercial office building located in Japan [1]. The proposed and baseline models are trained using the data from 2021-2022, and the simulations use data from 2023. We simulated the proposed algorithms and baselines for four different representative weeks during each of the seasons in 2023: Jan 23-Jan 30 (Winter), April 24-April 30 (Spring), Jul 17-Jul 23 (Summer), Oct 23-Oct 30 (Fall). The details on the disturbance data and simulation setup are given in Appendix F.

	Winter / Jan 2023		Spring / April 2023		Summer / Jul 2023		Fall / Oct 2023	
Algorithm	TD	Cost	TD	Cost	TD	Cost	TD	Cost
Gaussian	3.26	-16.01	9.66	6.02	21.14	19.48	16.86	14.60
VAR	1.04	-16.90	1.73	-2.39	2.03	10.15	2.90	6.02
Cantelli	0.10	-13.98	0.02	-1.97	0.06	9.78	0.17	5.16
Quantile	<u>0.47</u>	-17.88	<u>0.43</u>	-3.33	<u>0.66</u>	8.97	<u>0.36</u>	3.49

Table 1: Simulation results and baseline algorithms across seasonal test weeks. The best in each column is set in **bold** and the second best is underlined. Thermal discomfort (TD) is in units of deg-C·hr and cost is in scaled currency units. A negative cost implies no grid interaction; i.e. self-sustained operation.

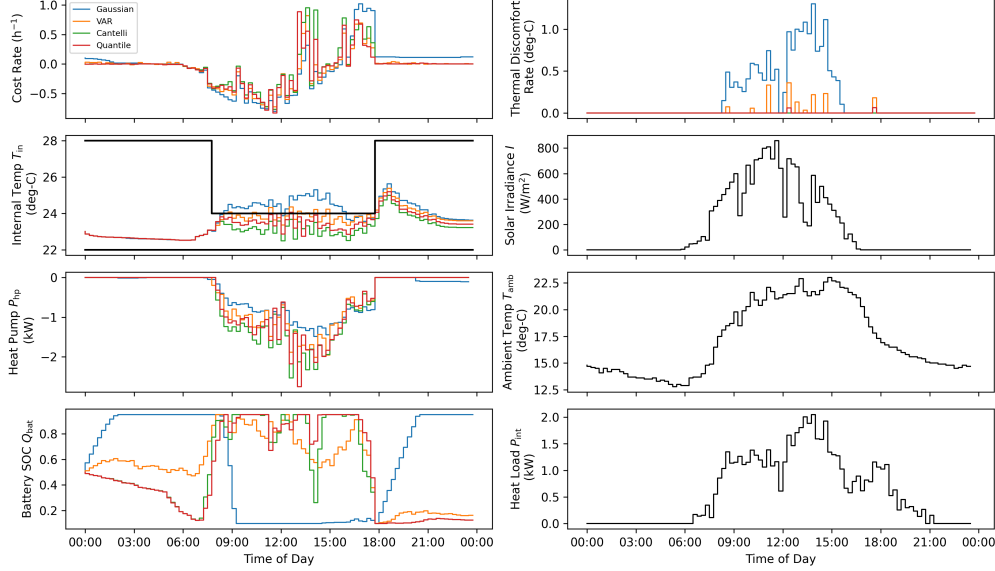


Figure 1: Trajectories of the proposed and baseline algorithms for an example day (Oct 2023). The temperature constraints are shown as a black line on the internal temperature T_{in} plot.

Simulation Results. The simulation results are shown in Table 1. These results are reported in terms of the net cost of electricity bought and sold from the grid (in scaled units) where negative cost indicates a net profit; and thermal discomfort (TD): the cumulative violation of the comfort constraints in units of Celsius-hours. In comparison to the baselines, our proposed algorithm (labeled Quantile) is either first or second for every metric and week tested. This balancing of cost and thermal discomfort was not matched by any of the baselines.

We also show the trajectories of the algorithms and the disturbances for an example day in Figure 1. This plot illustrates how Cantelli stays farther from the constraint boundary than Quantile, at the expense of greater heat pump use and electricity costs. This is despite the fact that both Cantelli and Quantile are using the same forecasting model, confirming that the Cantelli approach is a more conservative way to estimate the chance-constraints from forecast samples. Furthermore, Figure 1 also shows that Gaussian and VAR have a significant amount of constraint violation, which is consistent with the results in Table 1. This is likely due to the Gaussian and VAR disturbance models being unable to capture the true distribution of the disturbances.

Across all seasons (see Appendix G for supporting empirical evidence), we observe a clear trade-off between thermal comfort and operational cost, with uncertainty-aware methods consistently outperforming naive baselines. Gaussian assumptions lead to frequent violations of comfort bounds (up to 21 deg-C·hr in summer) as they underestimate forecast uncertainty and overdraw storage, while VAR achieves low cost but at the expense of sustained discomfort. In contrast, Cantelli’s inequality-based controller nearly eliminates discomfort (≤ 0.2 deg-C·hr across seasons) while maintaining competitive cost, reflecting a conservative but safe strategy that leverages thermal inertia and avoids aggressive battery cycling. Our proposed quantile method provides the best overall balance, delivering

the lowest cost in three of four seasons while keeping discomfort low (0.4–0.7 deg-C·hr), effectively timing heat pump operation and battery discharge to align with solar availability and peak loads. Single-day trajectories confirm these patterns: robust controllers act proactively by pre-cooling or pre-heating during off-peak hours, and reserving battery charge for midday stress periods, whereas Gaussian control depletes flexibility too early and reacts too late. Together, these results highlight that our risk-aware optimization provides a principled way to exploit building thermal mass and on-site storage, simultaneously safeguarding comfort and shaving operational cost.

References

- [1] SUSTIE is Mitsubishi Electric’s net Zero Energy Building test facility. URL <https://www.mitsubishielectric.com/rd/sustie/index.html>. (Available online).
- [2] American Society of Heating, Refrigerating and Air-Conditioning Engineers. Nonresidential cooling and heating load calculations. In *ASHRAE Handbook – Fundamentals*, chapter 18. ASHRAE, Peachtree Corners, GA, 2025. Table 1: Heat Gain from Occupants of Conditioned Spaces.
- [3] G. C. Calafiore and M. C. Campi. The scenario approach to robust control design. *IEEE Transactions on Automatic Control*, 51(5):742–753, 2006. doi: 10.1109/TAC.2006.875041.
- [4] J. Drgoňa, J. Arroyo, D. Vrabie, L. Helsen, and M. Wetter. All you need to know about model predictive control for buildings. *Annual Reviews in Control*, 50:190–232, 2020.
- [5] M. Farina, L. Giulioni, L. Magni, and R. Scattolini. A probabilistic approach to model predictive control. In *52nd IEEE conference on decision and control*, pages 7734–7739. IEEE, 2013.
- [6] N. He, J. Guo, Y. Li, Y. Quan, R. Li, and L. Yang. Stochastic model predictive control for the optimal operation of office buildings. *Building and Environment*, 267:112248, 2025. doi: 10.1016/j.buildenv.2024.112248.
- [7] K. Kan, F.-X. Aubet, T. Januschowski, Y. Park, K. Benidis, L. Ruthotto, and J. Gasthaus. Multivariate quantile function forecaster. In *International Conference on Artificial Intelligence and Statistics*, pages 10603–10621. PMLR, 2022.
- [8] B. Kouvaritakis, M. Cannon, S. V. Raković, and Q. Cheng. Explicit use of probabilistic distributions in linear predictive control. *Automatica*, 46(10):1719–1724, 2010.
- [9] M. Langner, J. Umlauf, C. Brosinsky, T. Leibfried, and A. Fuchs. Hierarchical-stochastic model predictive control (h-smpc) for the optimal energy management of campus buildings. *Journal of Building Engineering*, 89:109401, 2024. doi: 10.1016/j.jobbe.2024.109401.
- [10] B. Lim, S. Ö. Arık, N. Loeff, and T. Pfister. Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International journal of forecasting*, 37(4):1748–1764, 2021.
- [11] H. Lütkepohl. *New Introduction to Multiple Time Series Analysis*. Springer, 2005.
- [12] A. Mesbah. Stochastic model predictive control: An overview and perspectives for future research. *IEEE Control Systems Magazine*, 36(6):30–44, 2016. doi: 10.1109/MCS.2016.2602087.
- [13] P. Mohebi, S. Li, and Z. Wang. Chance-constrained stochastic framework for building thermal control under forecast uncertainties. *Energy and Buildings*, 331:115385, 2025.
- [14] F. Oldewurtel, A. Parisio, C. N. Jones, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, and M. Morari. Use of model predictive control and weather forecasts for energy efficient building climate control. *Energy and Buildings*, 45:15–27, 2012.
- [15] F. Oldewurtel, C. N. Jones, A. Parisio, and M. Morari. Stochastic model predictive control for energy efficient building climate control. *IEEE Transactions on Control Systems Technology*, 21(6):2618–2627, 2013.

- [16] J. A. Paulson, E. A. Buehler, R. D. Braatz, and A. Mesbah. Stochastic model predictive control with joint chance constraints. *International Journal of Control*, 93(1):126–139, 2020.
- [17] S. S. Rangapuram, L. Stella, D. Arbour, A. Smola, Jan Gasthaus, and T. Januschowski. Deep state space models for time series forecasting. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [18] K. Rasul, C. Seward, I. Schuster, and R. Vollgraf. Autoregressive denoising diffusion models for multivariate probabilistic time series forecasting. *arXiv preprint arXiv:2101.12072*, 2021.
- [19] K. Rasul, A.-S. Sheikh, I. Schuster, U. Bergmann, and R. Vollgraf. Multivariate probabilistic time series forecasting via conditioned normalizing flows. In *International Conference on Learning Representations*, 2021.
- [20] L. Romao, A. Papachristodoulou, and K. Margellos. On the exact feasibility of convex scenario programs with discarded constraints. *IEEE Transactions on Automatic Control*, 68(4):1986–2001, 2022.
- [21] D. Salinas, M. Bohlke-Schneider, L. Callot, R. Medico, and J. Gasthaus. High-dimensional multivariate forecasting with low-rank gaussian copula processes. In *Advances in Neural Information Processing Systems*, 2019.
- [22] D. Salinas, V. Flunkert, J. Gasthaus, and T. Januschowski. DeepAR: Probabilistic forecasting with autoregressive recurrent networks. *International Journal of Forecasting*, 36(3):1181–1191, 2020. doi: 10.1016/j.ijforecast.2019.07.001.
- [23] A. Satchwell, M. A. Piette, A. Khandekar, J. Granderson, N. M. Frick, R. Hledik, A. Faruqui, L. Lam, et al. A national roadmap for grid-interactive efficient buildings. Technical report, U.S. Department of Energy, Lawrence Berkeley National Laboratory and The Brattle Group, 2021. May 2021.
- [24] G. Schildbach, L. Fagiano, C. Frei, and M. Morari. The scenario approach for stochastic model predictive control with bounds on closed-loop constraint violations. *Automatica*, 50(12): 3009–3018, 2014.
- [25] F. Soroufifar, J. A. Paulson, Y. Wang, R. Quirynen, C. R. Laughman, and A. Chakrabarty. Bayesian forecasting with deep generative disturbance models in stochastic MPC for building energy systems. In *2024 IEEE Conference on Control Technology and Applications (CCTA)*, pages 414–419. IEEE, 2024.
- [26] W.-T. Tang, A. P. Vinod, F. G. Germain, and A. C. Joel A. Paulson, Christopher R. Laughman. Ai-driven scenario discovery: Diffusion models and multi-armed bandits for building control validation. Manuscript submitted for publication, 2025.
- [27] A. P. Vinod and S. Di Cairano. Sample quantile-based programming for non-convex separable chance constraints. In *2023 American Control Conference (ACC)*, pages 1517–1522. IEEE, 2023.
- [28] Y. Yang, G. Hu, and C. J. Spanos. Stochastic optimal control of HVAC system for energy-efficient buildings. *IEEE Transactions on Control Systems Technology*, 30(1):376–383, 2021.
- [29] X. Yuan and Y. Qiao. Diffusion-ts: Interpretable diffusion for general time series generation. In *International Conference on Learning Representations*, 2024.
- [30] H. Zhou, S. Zhang, J. Peng, S. Zhang, J. Li, H. Xiong, and W. Zhang. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 11106–11115, 2021.
- [31] T. Zhou, Z. Ma, Q. Wen, X. Wang, L. Sun, and R. Jin. Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. In *International conference on machine learning*, pages 27268–27286. PMLR, 2022.

A Related Work

A.1 Probabilistic time-series forecasting

Modern multi-horizon forecasting has shifted from point estimates to calibrated predictive distributions, with deep sequence models now standard. Temporal Fusion Transformers (TFT) unify attention with gating for interpretable, high-accuracy multi-horizon forecasts and natively support quantile training [10]. Autoregressive recurrent approaches such as DeepAR established scalable likelihood-based density forecasting on large related series [22], while deep state-space models learn latent probabilistic dynamics amenable to principled uncertainty propagation [17]. Transformer variants optimized for long contexts improve efficiency for extended horizons; e.g. Informer [30] and FedFormer [31] report good results on energy system-relevant time-series.

Most relevant for our work, multi-variate probabilistic forecasters aim to model the full joint uncertainty across variables rather than predicting each series in isolation. Import multivariate architectures include: copula-based models [21], which uses a Gaussian copula to represent dependencies between variables; flow-based models [19], which uses conditional normalizing flow trained with negative log likelihood minimization; and diffusion-based models [18, 29], which generate future paths by iteratively denoising at each step. Our proposed multi-variate forecasting model builds on the Multivariate Quantile Function Forecaster (MQF2) architecture [10], which handles multivariate distributions by using the gradient of an input-convex neural network to represent the multi-variate quantile function and is trained via energy score (although it can also be trained with conditional flows). We use the MQF2 as a distribution head on top of a TFT, which allows for multi-variate probabilistic forecasting while maintaining the benefits of the TFT architecture, such as robust handling of heterogeneous covariates and step-wise variable selection.

A.2 Stochastic MPC with disturbance forecasting

SMPC is a version of the classical MPC algorithm that explicitly models the impact of random disturbances on the system state when solving the constrained optimal control problem. Typically, constraint satisfaction is enforced with high probability with respect to possible state realizations, i.e. via chance-constraints. Most works within this vast literature assume that the disturbances are Gaussian [12], which is often insufficient to accurately model the disturbances in BEMS. One popular approach to circumvent the assumption of Gaussianity is to use the Cantelli inequality to construct conservative bounds on the chance-constraint with only the mean and variance of the disturbance distribution, e.g. [8, 5, 16]. This is different from the empirical quantile-based SMPC approach that we adopt from [27], which directly estimates the chance-constraints through samples and guarantees high-probability satisfaction of chance-constraints for arbitrary distributions. In our BEMS simulations, we find that the quantile-based approach more effectively balances constraint satisfaction and cost minimization than the Cantelli inequality-based approach.

Another direction for handling non-Gaussian disturbances in SMPC is the so-called scenario approach, in which several disturbance sequences are sampled and then the constraint are enforced with respect to all of these sequences or a selected subset of them, e.g. [3, 24, 20, 25]. Although the scenario approach guarantees constraint satisfaction with high probability after so many samples, it is often conservative with respect to the constraints [12]. It can also be computationally expensive, as it requires a duplicate set of optimization constraints for every sampled trajectory, versus the quantile-based SMPC that we use, which only requires the original constraints. We leave an empirical comparison between quantile and scenario-based SMPC in BEMS as future work.

A.3 GEB system management and control with SMPC

GEBs coordinate flexible end uses, energy storage, and on-site generation to provide both building-level performance and grid services. The U.S. DOE’s national roadmap frames GEBs as a controls problem at scale, emphasizing load flexibility, interoperability, and valuation of services [23]. At the controller level, MPC/SMPC has been repeatedly validated for HVAC energy efficiency and comfort, and extended to grid services. Early building SMPC studies quantified efficiency gains and comfort-risk trade-offs using weather forecasts [14], while follow-ons targeted tractable formulations for large systems and probabilistic constraints [15]. Aggregations of commercial buildings have been shown capable of reliable frequency-reserve provision under robust/SMPC formulations with energy

constraints [9, 6]. Comprehensive surveys document MPC/SMPC adoption across buildings, storage, and DER coordination, and summarize modeling and deployment hurdles (identification, uncertainty modeling, computational tractability, and KPI alignment) [4].

B GEB Model Description




Figure 2: Illustration of building model.

Table 2: Building model parameters.

Parameter	Value	Parameter	Value
$C_{w_{ext}}$	1.18×10^7	$C_{w_{int}}$	1.18×10^7
C_{in}	6.66×10^5	C_{itm}	1.63×10^7
R_{wall}	45×10^{-3}	R_{roof}	54×10^{-2}
R_{floor}	∞	R_{wind}	1×10^{-1}
R_{door}	19×10^{-2}	$R_{w_{int}}$	71×10^{-4}
R_i	22×10^{-4}	f	0.30
η_{heat}	3.50	η_{cool}	2.70
η_{sol}	0.19	η_{ch}, η_{dis}	0.93×10^{-6}
$F_{w_{ext} \rightarrow w_{in}}$	0.50	$F_{roof \rightarrow in}$	0.20
$F_{w_{in} \rightarrow in}$	0.30	$\alpha_{w_{ext}}$	0.13
α_{roof}	0.05	h_{roof}	3.45
$h_{w_{ext}}$	20.0	$A_{w_{ext}}$	52.2
A_{wind}	6.10	A_{roof}	65.9
N_{pv}	25	A_{pv}	1.685 sq. m.
N_{oct}	50	T_{stc}	25.0 deg-C
η_{stc}	0.19	$\eta_{l\&t}$	0.90
η_T	0.005	T_{noc}	25.0 deg-C

We use the building model from [26], which we restate in the following. An illustration of this model is shown in Figure 2. The model parameter values are shown in Table 2.

B.1 5-State Model of Thermal Dynamics

The thermal model consists of a single thermal zone and represent the envelope with a four-node RC network. The state collects the exterior and interior wall temperatures ($T_{w_{ext}}, T_{w_{int}}$), the zone air

temperature T_{in} , and the lumped internal thermal mass T_{itm} . The manipulated input is the heat pump's delivered heating/cooling rate P_{hp} . Exogenous signals are (i) ambient air temperature T_{amb} , (ii) shortwave irradiance I that drives solar gains, and (iii) internally generated heat P_{int} from equipment and occupants. The resulting continuous-time dynamics are

$$C_{w_{ext}} \dot{T}_{w_{ext}} = \frac{3}{R_{wall}} (T_{saw} + T_{w_{int}} - 2T_{w_{ext}}) \quad (5a)$$

$$C_{w_{int}} \dot{T}_{w_{int}} = \frac{3}{R_{wall}} (T_{w_{ext}} + T_{sin} - 2T_{w_{int}}) \quad (5b)$$

$$C_{in} \dot{T}_{in} = \frac{T_{amb} - T_{in}}{R_{wind}} + \frac{T_{amb} - T_{in}}{R_{door}} + \frac{T_{sin} - T_{in}}{R_{w_{int}}} + \frac{T_{itm} - T_{in}}{R_i} + \frac{T_{sar} - T_{in}}{R_{roof}} - P_{hp} + P_{int} + P_{sol_{in}} \quad (5c)$$

$$C_{itm} \dot{T}_{itm} = \frac{T_{in} - T_{itm}}{R_i} + \frac{T_{grnd} - T_{itm}}{R_{floor}} + P_{sol_{itm}}, \quad (5d)$$

where $C_{w_{ext}}$, $C_{w_{int}}$, C_{in} , C_{itm} are thermal capacitances and R_{door} , R_{wall} , $R_{w_{int}}$, R_{wind} , R_i , R_{floor} , R_{roof} are the corresponding resistances. When the floor is effectively insulated we take $R_{floor} = \infty$, which removes the dependence on ground temperature T_{grnd} . The inner wall surface temperature T_{sin} balances conduction through the wall and exchange with the zone air, while effective outer-surface temperatures T_{saw} and T_{sar} account for the effect at the exterior wall and roof, respectively. These relate to the other variables as,

$$(T_{w_{int}} - T_{sin})/(R_{wall}/3) = (T_{sin} - T_{in})/R_{w_{int}}, \quad (6a)$$

$$T_{saw} = (\alpha_{w_{ext}})/(h_{w_{ext}}) F_{w_{ext}} I + T_{amb}, \quad (6b)$$

$$T_{sar} = (\alpha_{w_{roof}})/(h_{w_{roof}}) F_{w_{roof}} I + T_{amb}. \quad (6c)$$

The impact of the solar radiation on the space is modeled as,

$$P_{sol_{wind}} = A_{wind} F_{wind} \eta_{sol} I, \quad (7)$$

which we partition between the air and internal mass via a fixed fraction $f \in [0, 1]$:

$$P_{sol_{in}} = f P_{sol_{wind}}, \quad P_{sol_{itm}} = (1 - f) P_{sol_{wind}}. \quad (8)$$

B.2 Battery/PV subsystem

In addition to thermal actuation, the site can exchange electrical energy using a battery and PV array. We model the battery state-of-charge proxy Q_{bat} with integral charge/discharge dynamics,

$$\dot{Q}_{bat} = \eta_{bat_{chg}} P_{bat_{chg}} - \frac{P_{bat_{dis}}}{\eta_{bat_{dis}}}, \quad (9)$$

subject to mutually exclusive operating modes at each instant,

$$P_{bat_{chg}} = 0 \text{ or } P_{bat_{dis}} = 0, \quad (10)$$

with distinct charge/discharge efficiencies $\eta_{bat_{chg}}$, $\eta_{bat_{dis}}$.

B.3 Discretization.

Using zero-order hold with sampling time Δt , a linear, time-invariant, discrete-time model of the augmented state is obtained, i.e.,

$$x_{t+1} = Ax_t + Bu_t + Fw_t, \quad (11)$$

where $x_t = [T_{w_{ext}}, T_{w_{int}}, T_{in}, T_{itm}, Q_{bat}]$, $u_t = [P_{hp}, P_{bat_{chg}}, P_{bat_{dis}}]$, and $w_t = [T_{amb}, I, P_{int}]$. The matrices A, B, F follow from (5)–(9).

B.4 Grid interaction

Let P_{buy} denote the net grid exchange at each time step,

$$P_{\text{buy}} = \underbrace{(P_{\text{bat}_{\text{chg}}} + |P_{\text{hp}}|)}_{\text{power consumed}} - \underbrace{(P_{\text{bat}_{\text{dis}}} + P_{\text{pv}})}_{\text{power produced}}. \quad (12)$$

Here, P_{pv} is the instantaneous power produced by the photovoltaic (PV) array. We represent P_{pv} as

$$P_{\text{pv}} = N_{\text{pv}} A_{\text{pv}} \eta_{\text{stc}} \eta_{\text{l\&t}} I \left[1 - \left(\eta_{\text{T}} \left(T_{\text{amb}} + \frac{T_{\text{noc}} - 20}{800} I - T_{\text{stc}} \right) \right) \right], \quad (13)$$

where N_{pv} is the number of panels, each with area A_{pv} ; η_{stc} is the module efficiency under Standard Test Conditions (STC) and $\eta_{\text{l\&t}}$ captures aggregate balance-of-system losses. The output scales with irradiance I and is de-rated by temperature effects inside the brackets, which depend on ambient temperature T_{amb} , the Nominal Operating Cell Temperature (NOCT) T_{noc} , the STC reference T_{stc} , and a temperature coefficient η_{T} . The NOCT-based term provides a simple proxy for cell temperature under non-STC irradiance.

C Forecasting Model Details

In the following, we describe the components of our proposed forecasting model, and provide an empirical comparison with vanilla TFT. The hyperparameters are given in Table 3. The model was trained on 12 GB GPUs.

Table 3: Hyperparameters for the proposed forecasting model.

Component	Hyperparameter	Value
Windowing	look-back (k)	768
	horizon (H)	96
AR-TFT	hidden size	64
	num lstm layers	1
	num attention heads	8
	dropout	0.2
PICNN (in MQF2)	latent dimension	3×13
	hidden size	32
	num layers	4
Training	optimizer	AdamW
	learning rate	1×10^{-4}
	batch size	16
	num epochs	10
	num energy score samples	50

C.1 TFT with Autoregressive Feedback (AR-TFT)

We use a modified version of the standard TFT architecture [10] that incorporates autoregressive feedback, which we call AR-TFT. In particular, for each horizon τ , the decoder input is both the covariate and a feedback term ($c_{t+\tau}, \bar{w}_{t+\tau-1}$), instead of just the covariate $c_{t+\tau}$ as in the standard TFT. For training, the feedback term $\bar{w}_{t+\tau-1}$ is taken to be the target value at the previous horizon (i.e. $\bar{w}_{t+\tau-1} = w_{t+\tau-1}$) and, for inference, the feedback term is taken to be the forecast sample at the previous horizon (i.e. $\bar{w}_{t+\tau-1} = \hat{w}_{t+\tau-1}$). This modified TFT architecture is used to output a horizon-specific latent vector, which we denote as,

$$\mathbf{h}_\tau = F_\theta(w_{t-k:t}, c_{t-k:t+\tau}, \bar{w}_{t+1:t+\tau-1}, \tau), \quad \tau = 1, \dots, H,$$

C.2 MQF2 Conditional Sampler

We use the MQF2 architecture [7] to produce a multivariate generative forecast from the latent vector at each horizon. The MQF2 sampler $g_\phi : (\mathbf{h}, \alpha) \mapsto \hat{w}$ is common for all horizons. The MQF2 architecture consists of a partially input-convex neural network (PICNN) denoted by $\Phi_\phi(\alpha; \mathbf{h})$ whose gradient with respect to α is used to map the latent and uniform random vector to the forecast,

$$\hat{w}_{t+\tau} = g_\phi(\mathbf{h}_\tau, \alpha) := \nabla_\alpha \Phi_\phi(\alpha; \mathbf{h}_\tau), \quad \alpha \sim \mathcal{U}([0, 1]^m), \quad (14)$$

where m is the dimension of the disturbance, i.e. $m = 3$ in this setup.

C.3 Training via Energy Score

We train end-to-end by minimizing the multivariate energy score. In particular, denote the true disturbance sequence as $\mathbf{w} = [w_{t+1}, \dots, w_{t+N}]$ and samples of the disturbance sequence forecast as $\hat{\mathbf{w}}^\ell = [\hat{w}_{t+1}^\ell, \dots, \hat{w}_{t+N}^\ell]$ for sample indices $\ell \in [M]$. Then, a Monte-Carlo estimate of the energy score is,

$$\widehat{\text{ES}} = \frac{1}{M} \sum_{\ell=1}^M \|\hat{\mathbf{w}}^\ell - \mathbf{w}\|_2 - \frac{1}{2M(M-1)} \sum_{\substack{\ell, \ell'=1 \\ \ell \neq \ell'}}^M \|\hat{\mathbf{w}}^\ell - \mathbf{w}^{\ell'}\|_2. \quad (15)$$

During training, we keep the model from the epoch with the best validation loss.

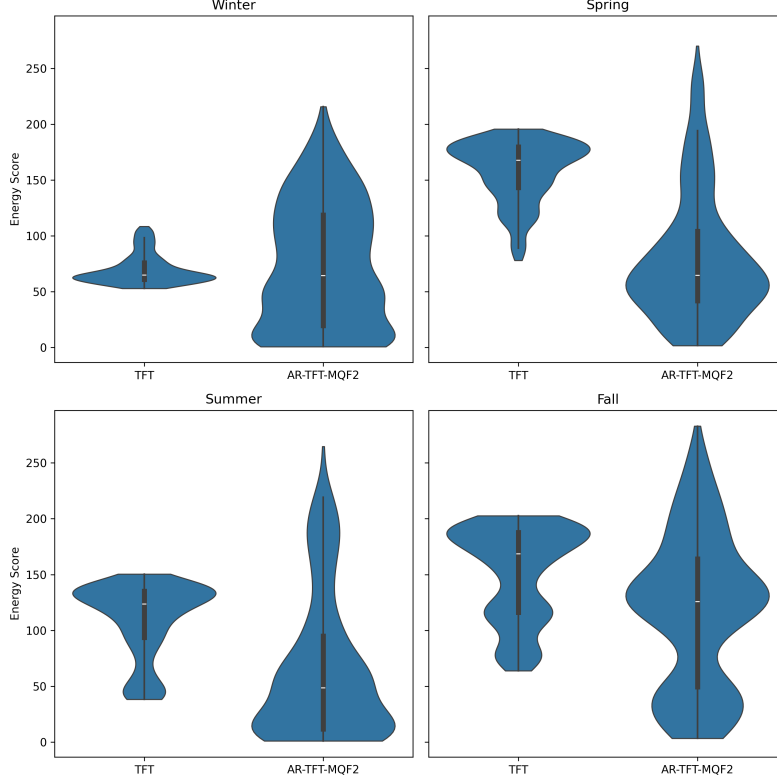


Figure 3: Comparison AR-TFT-MQF2 and vanilla TFT in terms of the energy score metric (lower is better) for seasonal test weeks.

Table 4: Hyperparameters of the vanilla TFT baseline.

Component	Hyperparameter	Value
Windowing	Look-back length (k)	768
	Forecast horizon (H)	96
TFT	Hidden size	32
	LSTM layers	2
	Attention heads	8
	Dropout	0.2
	Quantiles	0.01, 0.05, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 0.95, 0.99
Training	Optimizer	AdamW
	Learning rate	1×10^{-4}
	Batch size	16
	Epochs	10

C.4 Empirical Comparison with vanilla TFT

In this section, we give an empirical comparison between our proposed model (AR-TFT-MQF2) and the vanilla TFT model [10]. The hyperparameters for the vanilla TFT are given in Table 4. Note that the size of the output of the vanilla TFT (i.e. number of quantiles) is the same as the size of the output of the TFT module in our proposed model (13 per horizon and disturbance variable). In order to compare the two models at generative multi-variate forecasting, samples are generated from the univariate quantiles of the vanilla TFT model using quantile function sampling (see `QuantileRegression.sample()` in the `darts` documentation for details).

We compare the models in terms of the energy score metric (15) for four seasonal test weeks in 2023: Jan 23-Jan 30 (Winter), April 24-April 30 (Spring), Jul 17-Jul 23 (Summer), Oct 23-Oct 30 (Fall). The results are shown in Figure 3. Figure 3 shows seasonal violin plots of energy scores for the baseline TFT and the proposed AR-TFT-MQF2 model. In all seasons, AR-TFT-MQF2 attains lower medians, indicating sharper and better-calibrated forecasts. The performance gap is most pronounced in spring and summer, when disturbances such as solar irradiance, cooling demand, and occupancy schedules exhibit stronger diurnal and weekly variability as well as higher cross-variable correlations. While AR-TFT-MQF2 produces a wider spread of scores, the central tendency consistently outperforms TFT whose forecasts are narrower but biased due to its inability to capture inter-variable dependencies. These results confirm that auto-regressive MQF2 decoding enables TFT to learn both temporal and cross-variable correlations, yielding more faithful disturbance trajectories for downstream stochastic control.

D More Details on SMPC with Quantile Estimation

In this appendix, we provide additional details on the SMPC with conservative quantile estimation beyond what was provided in Section 2. In particular, Appendix D.1 describes in detail how the chance-constraints can be formulated with the empirical quantile function, and Appendix D.1 gives pseudocode for the proposed SMPC algorithm.

D.1 Quantile Formulation of Chance Constraint

In this section, we give details on how the SMPC chance constraints can be rewritten in terms of the quantile function. The original SMPC chance constraints are,

$$\begin{aligned} \mathbb{P}(C\hat{x}_\tau \leq T_{\text{in}}^{\max}) &\geq p \quad \forall \tau \in [t, t+N] \\ \mathbb{P}(C\hat{x}_\tau \geq T_{\text{in}}^{\min}) &\geq p \quad \forall \tau \in [t, t+N] \\ \hat{x}_{\tau+1} &= A\hat{x}_\tau + Bu_\tau + G\hat{w}_\tau \quad \forall \tau \in [t, t+N-1] \\ (\hat{w}_t, \dots, \hat{w}_{t+N-1}) &\sim f, \quad \hat{x}_\tau = x_t. \end{aligned} \tag{16}$$

Then, note that the predicted state at each time step can be written as,

$$\begin{aligned} \hat{x}_\tau &= A^{\tau-t}x_t + \sum_{s=t}^{\tau} A^{\tau-s}Bu_s + \sum_{s=t}^{\tau} A^{\tau-s}G\hat{w}_s \\ &= A^{\tau-t}x_t + [A^{\tau-t} \dots A I \mathbf{0} \dots \mathbf{0}]B\mathbf{u} + [A^{\tau-t} \dots A I \mathbf{0} \dots \mathbf{0}]G\hat{\mathbf{w}}, \end{aligned}$$

where $\mathbf{u} = [u_t^\top u_{t+1}^\top \dots u_\tau^\top]^\top$ and $\hat{\mathbf{w}} = [\hat{w}_t^\top \hat{w}_{t+1}^\top \dots \hat{w}_\tau^\top]^\top$. And therefore, $C\hat{x}_\tau$ can be written as,

$$\begin{aligned} C\hat{x}_\tau &= \underbrace{CA^{\tau-t}}_{S_\tau^x}x_t + \underbrace{C[A^{\tau-t} \dots A I \mathbf{0} \dots \mathbf{0}]B}_{S_\tau^u}\mathbf{u} + \underbrace{C[A^{\tau-t} \dots A I \mathbf{0} \dots \mathbf{0}]G}_{S_\tau^w}\hat{\mathbf{w}} \\ &= S_\tau^x x_t + S_\tau^u \mathbf{u} + S_\tau^w \hat{\mathbf{w}} \end{aligned}$$

It follows that that the chance-constraints in (16) can be written as,

$$\begin{aligned} \mathbb{P}_{\hat{\mathbf{w}} \sim f}(S_\tau^x x_t + S_\tau^u \mathbf{u} + S_\tau^w \hat{\mathbf{w}} \leq T_{\text{in}}^{\max}) &\geq p \quad \forall \tau \in [t, t+N], \\ \mathbb{P}_{\hat{\mathbf{w}} \sim f}(S_\tau^x x_t + S_\tau^u \mathbf{u} + S_\tau^w \hat{\mathbf{w}} \geq T_{\text{in}}^{\min}) &\geq p \quad \forall \tau \in [t, t+N]. \end{aligned}$$

Since the chance-constraint is separable with respect to the decision variable \mathbf{u} and the randomness $\hat{\mathbf{w}}$, this can then be equivalently defined in terms of the quantile function of $S_\tau^w \hat{\mathbf{w}}$ (denoted by $Q_{S_\tau^w \hat{\mathbf{w}}}(\cdot)$):

$$\begin{aligned} S_\tau^x x_t + S_\tau^u \mathbf{u} + Q_{S_\tau^w \hat{\mathbf{w}}}(p) &\leq T_{\text{in}}^{\max} \quad \forall \tau \in [t, t+N], \\ S_\tau^x x_t + S_\tau^u \mathbf{u} + Q_{S_\tau^w \hat{\mathbf{w}}}(1-p) &\geq T_{\text{in}}^{\min} \quad \forall \tau \in [t, t+N]. \end{aligned}$$

This can be approximated using the empirical quantile function $\hat{Q}_{S_\tau^w \hat{\mathbf{w}}}$ for an appropriately chosen probability threshold Δ :

$$\begin{aligned} S_\tau^x x_t + S_\tau^u \mathbf{u} + \hat{Q}_{S_\tau^w \hat{\mathbf{w}}}(\Delta) &\leq T_{\text{in}}^{\max} \quad \forall \tau \in [t, t+N], \\ S_\tau^x x_t + S_\tau^u \mathbf{u} + \hat{Q}_{S_\tau^w \hat{\mathbf{w}}}(1-\Delta) &\geq T_{\text{in}}^{\min} \quad \forall \tau \in [t, t+N]. \end{aligned} \tag{17}$$

We take $\Delta = p + \sqrt{\frac{\ln(1/\beta)}{2M}}$ for M samples, which [27] showed ensures that the constraints on \mathbf{u} in (17) are a sufficient condition for the original chance-constraints (16) with probability at least $1 - \beta$.

D.2 Complete Algorithm

The complete algorithm for our approach is shown in Algorithm 1. In addition to the aspects described in Section 2, the algorithm also includes the following:

- The cost function in the MPC optimization problem is the mean of the costs of the predicted trajectories for each of the forecast samples.
- The comfort constraints are relaxed with slack variables that are then penalized in the cost function. This ensures that the optimization problem remains feasible, even when chance-constraints cannot be satisfied.

Algorithm 1 SMPC with Conservative Quantile Estimation

input current state x_t , context \mathbf{c}_t , number of sampled scenarios M , validity threshold β , probability threshold p , slack variable penalty λ_T .

1: Draw samples from TFT model with context:

$$\hat{\mathbf{w}}^\ell = [\hat{w}_{t+1}^\ell, \dots, \hat{w}_{t+N}^\ell] \stackrel{iid}{\sim} f(\mathbf{c}_t) \quad \forall \ell \in [M]$$

2: Map samples through dynamics:

$$v_\tau^\ell = S_\tau^w \hat{\mathbf{w}}^\ell \quad \forall \tau \in [t, t+N]$$

3: Compute conservative upper and lower quantiles of $\{v_\tau^\ell\}_{\ell=1}^N$:

$$\bar{\epsilon}_\tau = \hat{Q}\left(\Delta; \{v_\tau^\ell\}_{\ell=1}^N\right), \quad \underline{\epsilon}_\tau = \hat{Q}\left(1 - \Delta; \{v_\tau^\ell\}_{\ell=1}^N\right), \quad \Delta := p + \sqrt{\frac{\ln(1/\beta)}{2M}}$$

4: Solve Quantile-SMPC optimization problem:

$$\begin{aligned} \min_{\hat{\mathbf{x}}, \hat{\mathbf{u}}} \quad & \frac{1}{M} \sum_{\ell=1}^M J(\hat{\mathbf{x}}, \hat{\mathbf{u}}) + \lambda_T (\|\bar{\mathbf{s}}\|_1 + \|\underline{\mathbf{s}}\|_1) \\ \text{s.t.} \quad & \hat{x}_{\tau+1}^\ell = A\hat{x}_\tau^\ell + B\hat{u}_\tau + G\hat{w}_\tau^\ell \quad \forall \tau \in [t, T+N], \\ & S_\tau^x x_t + S_\tau^u \hat{\mathbf{u}} + \bar{\epsilon}_\tau \leq T_{\text{in}}^{\max} + \bar{s}_\tau \quad \forall \tau \in [t, T+N], \\ & S_\tau^x x_t + S_\tau^u \hat{\mathbf{u}} + \underline{\epsilon}_\tau \geq T_{\text{in}}^{\min} - \underline{s}_\tau \quad \forall \tau \in [t, T+N], \\ & \hat{u}_\tau \in \mathcal{U} \quad \forall \tau \in [t, T+N] \\ & \hat{x}_t^\ell = x_t \quad \forall \ell \in [M] \\ & \bar{\mathbf{s}}, \underline{\mathbf{s}} \geq 0, \end{aligned}$$

5: Choose first control action $u_t = \hat{u}_t$.

- The first control input computed by the MPC optimization problem (u_t) is then applied to the system.
- There are input constraints $\hat{u}_\tau \in \mathcal{U}$ in the MPC optimization problem, e.g. heat pump capacity, battery charge/discharge limits, state-of-charge limits.

E Baselines

All of the baselines solve the same SMPC optimization formulation to choose the next control action. The only differences is the model for the disturbance, which we detail in each of the following sections.

E.1 Gaussian

This algorithm computes the SMPC using an iid Gaussian model for the disturbances. The mean and covariance of this Gaussian is taken to be the empirical mean $\hat{\mu}$ and empirical covariance $\hat{\Sigma}$ of the training dataset. Then, given this disturbance model, the chance-constraints in the SMPC (3) are constructed exactly as,

$$\begin{aligned} S_\tau^x x_t + S_\tau^u \mathbf{u} + Q_{S_\tau^w \hat{\mathbf{w}}}(p) &\leq T_{\text{in}}^{\text{max}}, \\ S_\tau^x x_t + S_\tau^u \mathbf{u} + Q_{S_\tau^w \hat{\mathbf{w}}}(1-p) &\geq T_{\text{in}}^{\text{min}}, \end{aligned} \quad (18)$$

where $Q_{S_\tau^w \hat{\mathbf{w}}}(\cdot)$ is the quantile function, which can be computed exactly given that $S_\tau^w \hat{\mathbf{w}} \sim \mathcal{N}(S_\tau^w \hat{\mu}, S_\tau^w \hat{\Sigma} (S_\tau^w)^\top)$, where $\hat{\mu} = [\hat{\mu} \dots \hat{\mu}]^\top$ and $\hat{\Sigma} = \text{diag}(\hat{\Sigma}, \dots, \hat{\Sigma})$.

E.2 VAR

This algorithm computes the SMPC using a vector auto-regressive (VAR) model for the disturbances. The VAR model predicts the current disturbance as,

$$\hat{w}_t = M_1 w_{t-1} + \dots + M_k w_{t-k} + \epsilon_t \quad \epsilon_t \sim \mathcal{N}(\mu_\epsilon, \Sigma_\epsilon), \quad (19)$$

where M_1, \dots, M_k are coefficient matrices that are fit from data and k is the look-back window. In particular, we fit a VAR model to the training dataset using the VAR module from the package `statsmodels`.² We use same look-back window $k = 768$ also used by our proposed model. Under the VAR model, the stacked forecast $\hat{\mathbf{w}} = [\hat{w}_t^\top, \dots, \hat{w}_{t+H}^\top]$ is linear with respect to $\epsilon_t, \dots, \epsilon_{t+H}$, and therefore it is Gaussian. The means $\hat{\mu}_\tau$ of the disturbance forecasts \hat{w}_τ are computed recursively as,

$$\begin{aligned} \hat{\mu}_t &= M_1 w_{t-1} + \dots + M_k w_{t-k} + \mu_\epsilon, \\ \hat{\mu}_{t+1} &= M_1 \hat{\mu}_t + M_2 w_{t-1} + \dots + M_k w_{t-k+1} + \mu_\epsilon, \text{ and so on...} \end{aligned}$$

In order to write the covariance of the disturbance forecasts, we note that,

$$w_\tau - \hat{\mu}_\tau = \sum_{i=0}^{\tau-t-1} \Psi_i \epsilon_{\tau-i},$$

where Ψ_i are the MA coefficients of the model, see e.g. [11] Section 2.2. Therefore, the covariance $\hat{\Sigma}$ of the stacked forecast $\hat{\mathbf{w}}$ can be written as,

$$\hat{\Sigma} = L \Sigma_\epsilon L^\top,$$

where L is the lower block-Toeplitz matrix with $L_{j,i} = \Psi_{j-i}$ for all $j \geq i$ and 0 otherwise, and $\Sigma_\epsilon = \text{diag}(\Sigma_\epsilon, \dots, \Sigma_\epsilon)$. Then the chance-constraints are computed as in (18) using the specified mean and covariance of the stacked disturbance forecast.

E.3 Cantelli

This algorithm computes a conservative version of the exact SMPC using the Cantelli inequality as is often done in the SMPC literature, e.g. [5, 16]. In particular, it computes the samples $v_\tau^\ell = S_\tau^w \hat{\mathbf{w}}^\ell, \forall \tau \in [t, t+N], \forall \ell \in [M]$ using the forecast samples $\hat{\mathbf{w}}^\ell$ same as in Algorithm 1. The empirical mean $\hat{\mu}_\tau$ and empirical variance $\hat{\sigma}_\tau$ of these samples is then computed for each τ . The Cantelli inequality says that,

$$\mathbb{P}(S_\tau^w \mathbf{w} \geq \hat{\mu}_\tau + a) \leq \frac{\hat{\sigma}_\tau}{\hat{\sigma}_\tau + a^2} \quad \forall a \geq 0 \iff \mathbb{P}\left(S_\tau^w \mathbf{w} \geq \hat{\mu}_\tau + \hat{\sigma}_\tau \sqrt{\frac{1-\delta}{\delta}}\right) \leq \delta \quad \forall \delta > 0.$$

²https://www.statsmodels.org/dev/generated/statsmodels.tsa.vector_ar.var_model.VAR.html

This inequality is used to specify the SMPC constraints, i.e.,

$$\begin{aligned} S_\tau^x x_t + S_\tau^u \mathbf{u} + \hat{\mu}_\tau + \hat{\sigma}_\tau \sqrt{\frac{1-\delta}{\delta}} &\leq T_{\text{in}}^{\text{max}}, \\ S_\tau^x x_t + S_\tau^u \mathbf{u} + \hat{\mu}_\tau - \hat{\sigma}_\tau \sqrt{\frac{1-\delta}{\delta}} &\geq T_{\text{in}}^{\text{min}}. \end{aligned} \tag{20}$$

F Experimental Setup

F.1 Disturbance Data

Table 5: Summary statistics for disturbance dataset.

	Ambient Temperature	Solar Irradiance	Occupancy
mean	17.6182	167.8220	2.4353
std	8.1052	260.3170	3.8565
min	-3.3000	0.0000	0.0000
25%	11.2533	0.0000	0.0000
50%	17.9133	1.4000	0.0000
75%	24.3800	262.8000	4.4667
max	36.5800	1218.8667	21.8000

The disturbance data was collected from the Mitsubishi Electric SUSTIE [1] building during 2021-2024. This dataset includes ambient temperature, solar irradiance, and occupancy, with a 15 minute sampling period from Jan 9 2021 through Dec 31 2024. The summary statistics for this dataset is shown in Table 5. The first year of this dataset is used for training and the second year for validation as detailed in Table 6.

F.2 MPC Cost and Input Constraints

In the simulations, we use the MPC cost and input constraints from [26] for both the proposed method and baselines, which is described as follows.

Cost Function The cost function is as follows for a horizon of length N :

$$J(\bar{u}, \bar{w}) = \sum_{t=1}^N \max((P_{\text{buy}})_t, 0) + \lambda_{\text{hp}} \|(P_{\text{hp}})_{1:N}\|_1, \quad (21)$$

with user-chosen weights $\lambda_{\text{hp}} > 0$. The term $\max((P_{\text{buy}})_t, 0)$ charges only for net purchases (since $P_{\text{buy}} < 0$ when selling back to the grid), while $(P_{\text{hp}})_{1:N}$ denotes the control sequence over the horizon. An ℓ_1 penalty on P_{hp} encourages sparse (i.e., temporally concentrated) actuation.

Input Constraints There are hard bounds on the battery state by constraining the battery SOC Q_{bat} within prescribed limits:

$$Q_{\text{bat}}^{\min} \leq Q_{\text{bat}} \leq Q_{\text{bat}}^{\max}. \quad (22)$$

Control inputs are restricted as follows:

$$P_{\text{hp}}^{\min} \leq P_{\text{hp}} \leq P_{\text{hp}}^{\max}, \quad (23a)$$

$$0 \leq P_{\text{bat_chg}} \leq P_{\text{bat_chg}}^{\max} (1 - s_{P_{\text{bat_chg}}}), \quad (23b)$$

$$0 \leq P_{\text{bat_dis}} \leq P_{\text{bat_dis}}^{\max} s_{P_{\text{bat_chg}}}, \quad (23c)$$

$$0 \leq s_{P_{\text{bat_chg}}} \leq 1. \quad (23d)$$

When $s_{P_{\text{bat_chg}}} \in \{0, 1\}$, the pair of bounds (23b)–(23c) enforces the disjunctive condition (10) (pure charge at $s_{P_{\text{bat_chg}}} = 0$ or pure discharge at $s_{P_{\text{bat_chg}}} = 1$). For a convex formulation, binary requirement the binary requirement is relaxed to the continuous interval constraint (23d).

F.3 Simulation Setup

The simulations use the building model described in Appendix B. The building dataset provides the occupancy, but not the internal heat loads, which are required for the building model. Therefore, for the simulations, we convert the occupancy data to heat loads using the following approximation,

$$P_{\text{int}} = \sum_{i=1}^{N_{\text{occ}}} \max(q_i, 0), \quad q_i \sim \mathcal{N}(\mu_q, \sigma_q), \quad (24)$$

Table 6: Training and validation ranges with counts.

Split	Start date	End date	Number of data points
Training	2021-01-09	2022-01-08	35040
Validation	2022-01-09	2023-01-08	35040

Table 7: Simulation parameters.

	Parameter	Symbol	Value
Environment	Initial state	x_0	[23 23 23 23 0.5]
	Mean heat load per person (BTU/h)	μ_q	450
	Std heat load per person (BTU/h)	σ_q	67.5
	Max heat pump (kW)	P_{hp}^{\max}	4.2
	Min heat pump (kW)	P_{hp}^{\min}	-4.2
	Max state-of-charge	Q_{bat}^{\max}	0.95
	Min state-of-charge	Q_{bat}^{\min}	0.1
	Max internal temperature (deg-C)	T_{in}^{\max}	24 from 08:00 to 18:00, else 28
	Min internal temperature (deg-C)	T_{in}^{\min}	22
MPC	Constraint penalty weight	λ_T	1e4
	Heat pump penalty weight	λ_{hp}	0.5
	Probability threshold	p	0.9
Quantile	Number of forecast samples	M	500
	Validity threshold	β	0.1
Cantelli	Number of forecast samples	M	500
VAR	Look-back window	k	768

where N_{occ} is the occupancy, μ_q, σ_q represents the mean and standard deviation of the heat load per a person, which we choose according to ASHRAE typical values for office work [2].

The formula (24) is also used to convert the training dataset to be in terms of heat load P_{int} , which is then used for fitting of the Gaussian and VAR baselines. The proposed forecasting model is trained directly with the occupancy data, so (24) is also used for converting the samples from the forecasting model to heat load forecasts. The reason that we fit the Gaussian and VAR baselines to the approximate heat load data (and not the occupancy data) is because the SMPC with Gaussian and VAR disturbance models is only tractable if they predict in the space of the disturbance w_t ; see Appendix E.1 and Appendix E.2.

The parameters used for the simulations are shown in Table 7.

G Additional Plots

The full trajectories for each of the test weeks are given in Figure 4 (winter), Figure 5 (spring), Figure 6 (summer), Figure 7 (fall).

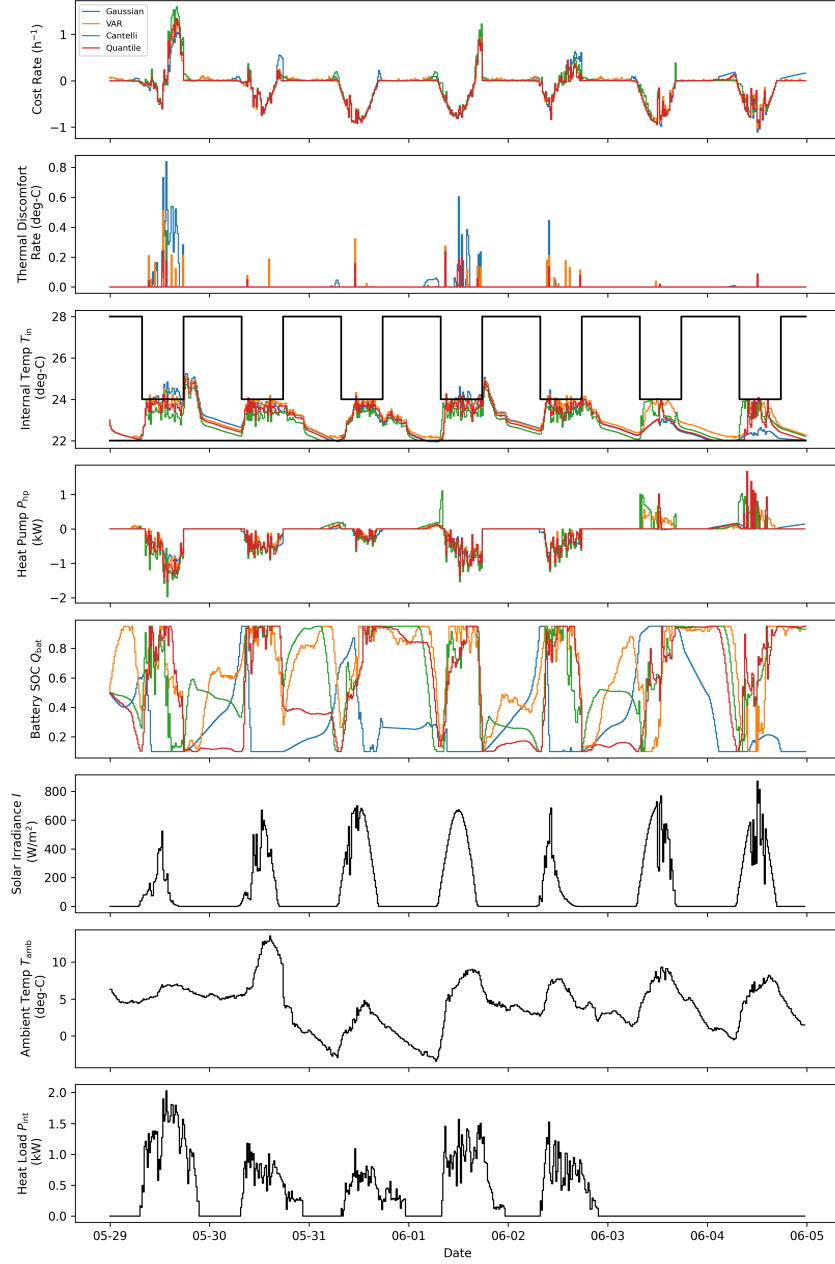


Figure 4: Trajectories of proposed and baseline SMPC algorithms during winter test week (January 23, 2023).

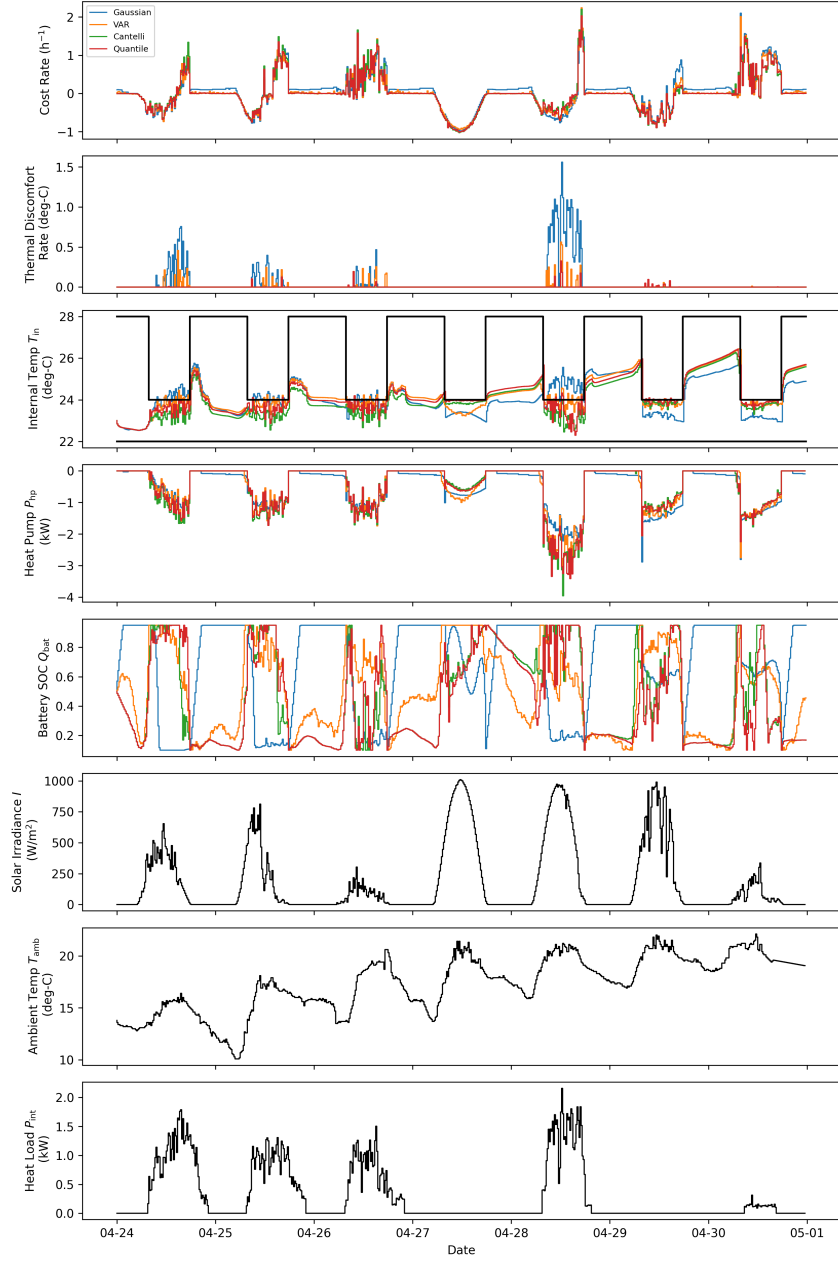


Figure 5: Trajectories of proposed and baseline SMPC algorithms during spring test week (April 24, 2023).

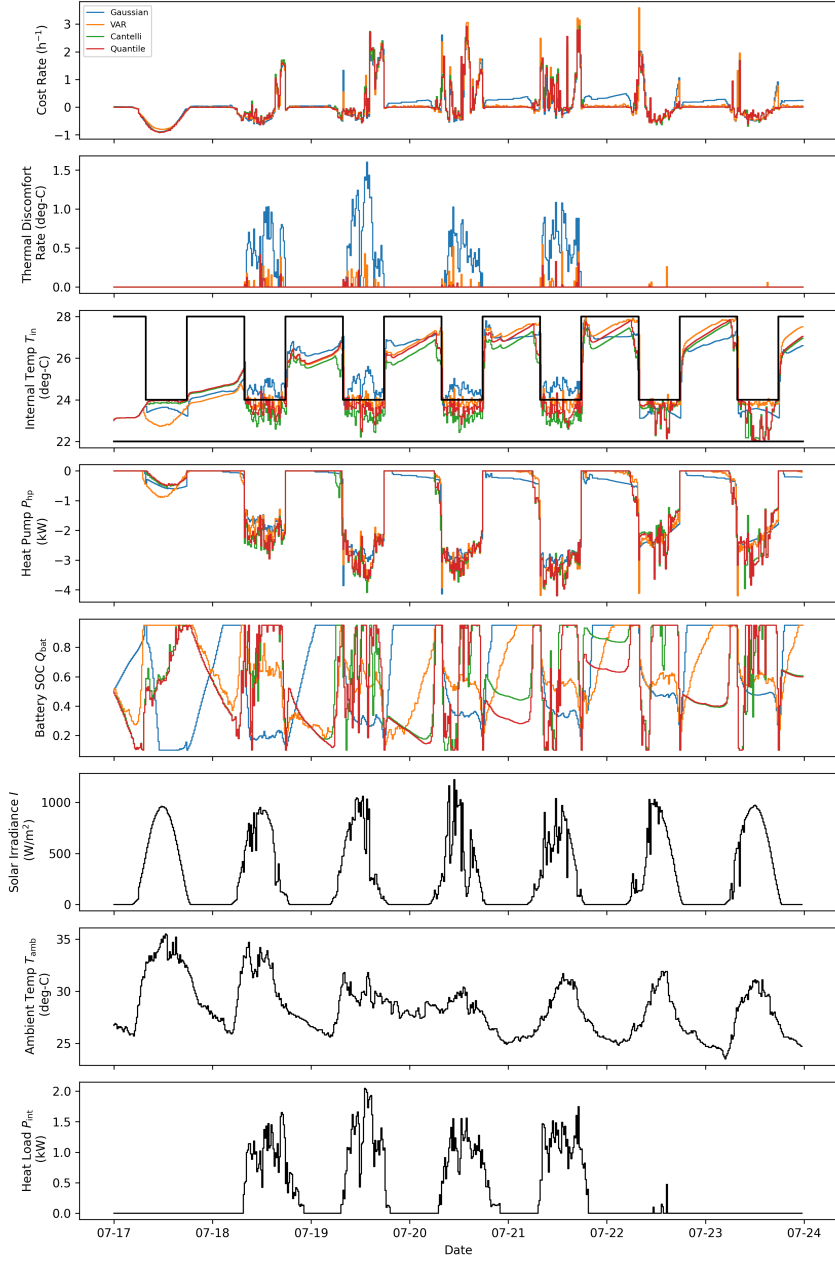


Figure 6: Trajectories of proposed and baseline SMPC algorithms during summer test week (July 17, 2023).

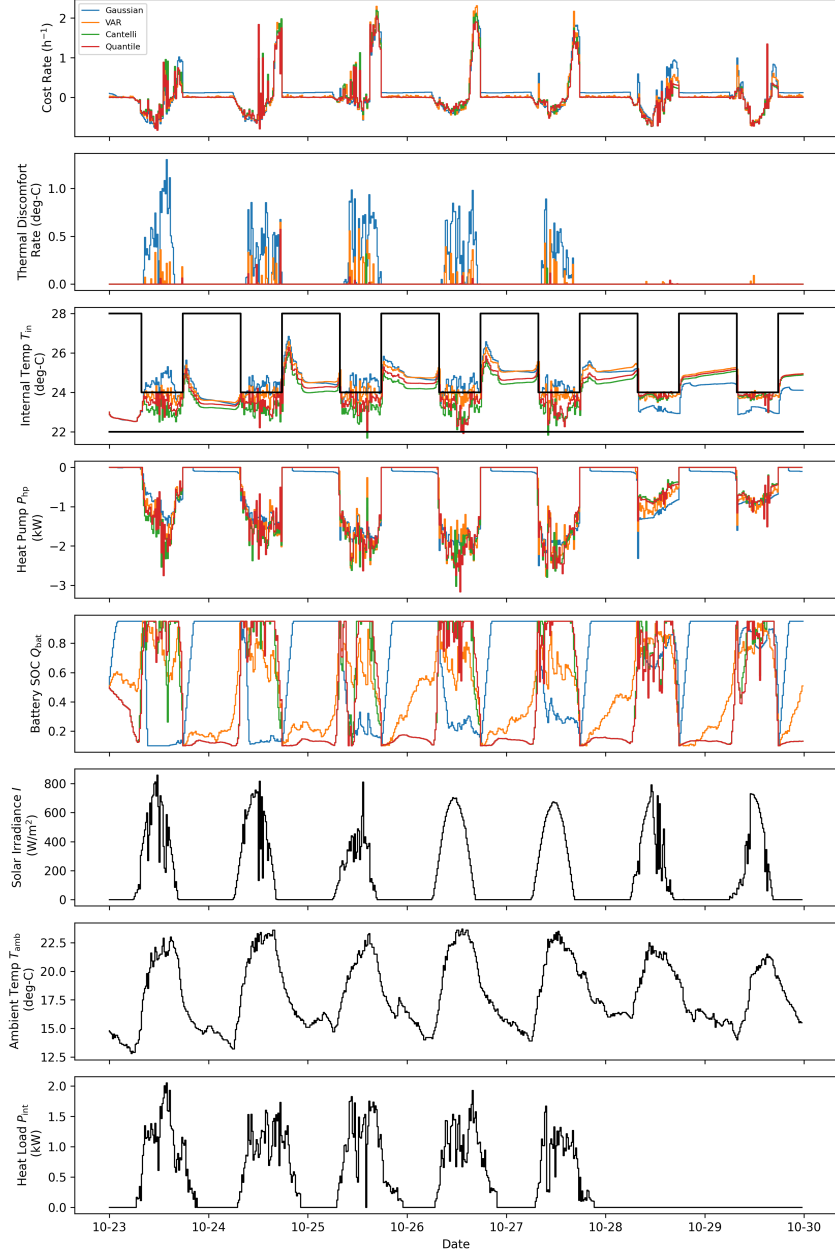


Figure 7: Trajectories of proposed and baseline SMPC algorithms during fall test week (October 23, 2023).