INL/JOU-25-83503-Revision-0

Expanding Market Opportunities: Cogeneration Strategies for Integrated PWR and Thermal Energy Storage Systems

November 2025

Changing the World's Energy Future

So-Bin Cho, Rami M Saeed, Xiaodong Sun, Todd Allen

Idaho National Laboratory

INL/JOU-25-83503-Revision-0

Expanding Market Opportunities: Cogeneration Strategies for Integrated PWR and Thermal Energy Storage Systems

So-Bin Cho, Rami M Saeed, Xiaodong Sun, Todd Allen

November 2025

Idaho National Laboratory
Idaho Falls, Idaho 83415

http://www.inl.gov

Prepared for the
U.S. Department of Energy
Under DOE Idaho Operations Office
Contract DE-AC07-05ID14517

Expanding Market Opportunities: Cogeneration Strategies for Integrated PWR and Thermal Energy Storage Systems

So-Bin Cho[1,2], Rami M. Saeed[2,*], Todd R. Allen[1], and Xiaodong Sun[1]

[1]Nuclear Engineering & Radiological Sciences, University of Michigan, Ann Arbor, MI, USA
[2]Idaho National Laboratory, Idaho Falls, ID, USA

ABSTRACT

We assess the economic viability of nuclear cogeneration by investigating three different modes—fixed dispatch, fully flexible dispatch, and flexible dispatch with minimum heat supply requirements. The analysis focuses on an existing pressurized water reactor (PWR) integrated with thermal energy storage (TES). Heat production costs are estimated under these modes for two U.S. electricity markets, the Electric Reliability Council of Texas (ERCOT) and the Pennsylvania–New Jersey–Maryland Interconnection (PJM). A sensitivity analysis is conducted at varying heat market prices to evaluate the profitability of PWR–TES cogeneration. Results indicate that fixed heat dispatch inflates heat production costs, often rendering projects economically feasible only at higher heat price levels. In contrast, fully flexible dispatch lowers heat production costs by an average of 43 % compared to fixed dispatch. However, the current 30% thermal dispatch limit may be insufficient to serve high baseline industrial demands cost-effectively; a higher maximum dispatch rate could enhance project economics. Additionally, while markets with higher and more volatile electricity prices (e.g., ERCOT) offer greater total energy sales potential (i.e., heat and electricity), they also increase opportunity costs if restrictions on heat production scheduling are required. In contrast, lower-price, less volatile markets (e.g., PJM) experience a smaller impact from such constraints and provide greater flexibility in accommodating varying cogeneration modes. Overall, these findings provide a framework to guide nuclear plant operators in aligning cogeneration strategies with industrial process requirements and electricity market conditions.

HIGHLIGHTS

- Heat demand and electricity price profiles guide optimal cogeneration mode selection.
- Flexible PWR cogeneration lowers costs by optimizing heat-electricity trade-offs.
- Cogeneration at a fixed rate increases heat production costs.
- Cogeneration in low electricity-price markets offers greater flexibility in heat supply.

KEYWORDS

energy arbitrage, nuclear cogeneration, opportunity cost, process heat market, thermal energy storage

INTRODUCTION

As of early 2025, the United States has 94 operating commercial nuclear reactors (the number recently increased with the startup of Vogtle Units 3 and 4 in 2023–4). These reactors are all light-water reactors, comprising 63 pressurized water reactors (PWRs) and 31 boiling water reactors (BWRs). They are spread across 55 plant sites in 28 states, supplying about 20% of the nation's total electricity and over half of its carbon-free generation. Since the shale gas revolution, the U.S. has effectively become self-sufficient and has even emerged as a net LNG exporter. This pushed Henry Hub prices down to sustained low levels (~\$2–3/MMBtu for much of 2015–2020), corresponding to \$7–\$10 per $MWh_{(t)}$ [1]. These low gas prices have been a double-edged sword: beneficial for consumers and gas-intensive industries but challenging

* Correspondence: rami.saeed@inl.gov

for nuclear power plants (NPPs). Notably, several U.S. nuclear plants were shut down largely due to market pressure from low gas prices [2].

Against this bleak backdrop, nuclear energy can offer significant benefits across the power and industrial sectors: it provides greenhouse gas (GHG)-free electricity and heat to non-grid applications by preventing the overbuilding of VRE and storage capacities [3, 4]. Studies also indicate that with tighter GHG constraints, nuclear energy becomes competitive, even at moderate cost levels [5]. Nuclear plants can achieve greater operational flexibility by cogeneration and/or coupling thermal energy storage (TES), allowing them to provide both baseload and peaking capacity [6]. Thus, expanding operational flexibility would enable nuclear energy to (1) provide GHG-free electricity or heat at a competitive price and (2) improve the economics, thereby extending the life of the existing U.S. reactor fleet.

Major initiatives driven in the 2000s and 2010s sought nuclear cogeneration on new reactor designs and direct coupling with process heat. The Nuclear Cogeneration Industrial Initiative (NC2I) [7] and Next Generation Nuclear Plant (NGNP) [8], launched under EU and USA leadership respectively, focused on high-temperature gas and hydrogen production. Each initiative yielded important outcomes: a conceptual HTGR cogeneration system design, a licensing framework, and a business plan. In parallel, today's cogeneration vision extends beyond standalone advanced reactor-heat projects to fully integrated energy systems. This often involves coupling existing nuclear reactors with other energy assets (renewables, storage, industrial facilities) to optimize overall efficiency and grid flexibility. The EU-funded NPHyCo (Nuclear Powered Hydrogen Cogeneration) is laser-focused on enabling large-scale hydrogen production utilizing existing nuclear power plants [9]. The Department of Energy (DOE)'s Light Water Reactor Sustainability (LWRS) Program [10] and Integrated Energy Systems (IES) [11] are more comprehensive efforts empowering nuclear energy's role beyond electricity generation.

The inclusion of existing PWR NPPs is particularly important because it demonstrates that cogeneration is broadly applicable across PWR-types advanced reactors (ARs), albeit with different reactor sizes. For instance, the Nine Mile Point NPP hydrogen project [12] shows how a BWR can dedicate a portion of its electrical output to a proton exchange membrane (PEM) electrolyzer, just as a PWR can. However, unlike BWRs, PWRs can additionally supply steam to a solid-oxide electrolyzer for higher efficiency [13]. Consequently, lessons from flexible LWRs can inform the design and operation of advanced reactors, potentially expediting flexibility and maximizing overall benefits.

Indeed, multiple studies under the LWRS [14-16] and IES programs [17-19] have explored PWR systems coupled to TES. Recent findings indicate that a generic four-loop Westinghouse PWR can support up to 30% heat dispatch with only minor adjustments to major components—such as turbines and condensers—and modest upgrades to ancillary equipment, including expansion joints and control valves [20]. Another study develops detailed cost models for TES integration in NuScale's VOYGR reactor, which is a small modular reactor (SMR) based on an integrated PWR design [17]. Despite this progress, the economic competitiveness of such flexibility has so far been evaluated predominantly in case-specific scenarios—such as electricity arbitrage or hydrogen production—and remains limited to particular electricity markets, steam-dispatch strategies, or industrial processes.

To address these gaps, this study proposes a mechanism-focused framework supported by parameterization and sensitivity analyses. This environment-agnostic approach aims to create broadly applicable insights for PWR cogeneration assessments. We model industrial applications as a heat market at different price points, enabling NPP owners to focus on the broader value of nuclear heat beyond any single industrial use. Furthermore, we segment this heat market customer by examining both a fixed 30% heat-dispatch mode and a flexible mode that allows dispatch up to 30% of the reactor's thermal capacity, thus capturing a spectrum of load types and operational constraints.

A two-tank molten-salt TES system is introduced to decouple the reactor's secondary loop from the heat market, thereby offering an additional physical barrier and buffering transients from the heat customer. Although various TES technologies could be employed, we focus on molten salt because it simplifies operational control [6, 10] and has a well-established commercialization pathway By incorporating the bottom-up capital cost estimates associated with TES integration, we provide a comparative evaluation of multiple PWR-TES integration scenarios against a business-as-usual (BAU) baseline.

METHODS

Building on the foundation established by previous studies [17, 20], we (1) parametrically formulate the PWR–TES system, (2) examine the optimal sizing of both TES and the balance of plant (BOP) as well as system-level economics measured by net present value (NPV), and (3) derive insights by relating key performance indicators—such as system capacity, heat production cost, and electricity or heat sales—through a comprehensive heat-price sensitivity analysis. We further compare the system's performance in two U.S. electricity markets, Electric Reliability Council of Texas (ERCOT) and Pennsylvania-New Jersey-Maryland Interconnection (PJM). Overall, the analysis is detailed enough to address specific industry needs while remaining sufficiently generalizable to provide actionable insights into the feasibility of PWR-TES integration. The following sections detail the development of the PWR–TES configurations and present the accompanying technoeconomic analysis (TEA) modeling strategies.

PWR-TES Coupling and Operation Strategies

This work considers a four-loop Westinghouse PWR rated at 3,650 $MW_{(t)}$, with a 30% heat-dispatch point (1,095 $MW_{(t)}$) located upstream of the high-pressure (HP) turbine. Under the BAU scenario, the plant operates in baseload mode exclusively for electricity production. To enhance flexibility while limiting changes to the existing PWR design, two coupling configurations are considered: (1) incorporating thermal energy storage (TES) alongside a secondary balance-of-plant (BOP) for electricity arbitrage and (2) directly supplying industrial consumers with stored heat. For non-electric uses, the secondary BOP can be replaced by an intermediate heat exchanger (IHX) connected to a single industrial heat user.

As illustrated in Figure 1, any heat stored in the TES is used exclusively for secondary BOP power generation or process-heat supply, and does not discharge back into the primary loop. While one could envision a single, oversized BOP that provides both electricity and process heat through a flexible TES loop, such a configuration would deviate from the baseline assumption of 30% heat dispatch and likely require significant additional investment and engineering analysis [20]. Accordingly, this study confines its scope to the nominal 30% dispatch point and a standard BOP capacity of 3,650 $MW_{(t)}$.

Under that constraint, two operating modes are explored: a fixed-dispatch mode, which may be employed for relatively constant heat demands (e.g., in an oil refinery), and a flexible-dispatch mode, which is suitable for industries with flexible heat demand scheduling (e.g., in a hydrogen plant). Additionally, different minimum-supply levels are introduced to reflect scenarios—such as industrial parks—that may have both a constant base-load component and a variable portion. By capturing these modes and associated load profiles, the analysis aims to capture a heterogenous process heat applications from a NPP owners point. By examining these operational modes, our model provides a parsimonious framework for exploring possible configurations and use cases at the early stages of PWR–TES integration, while capturing the key features associated with non-electricity applications.

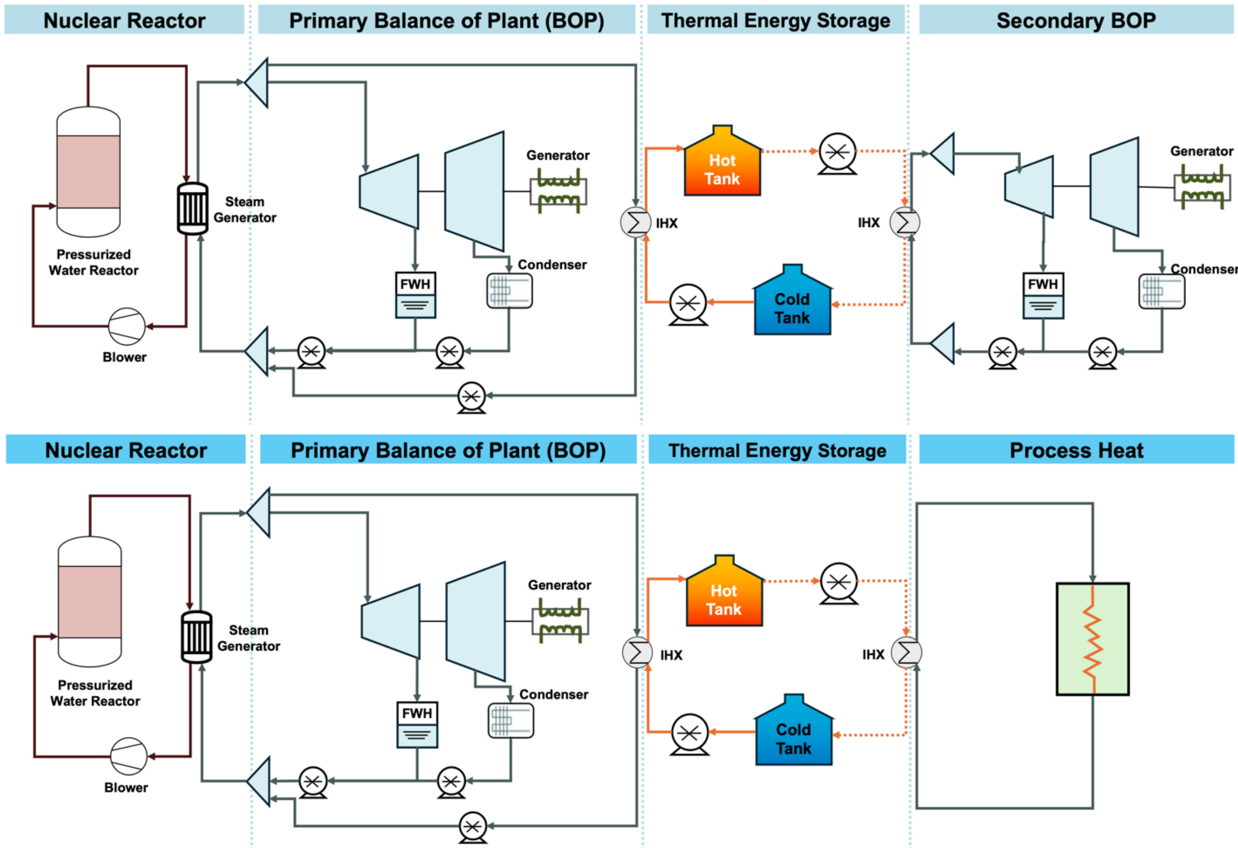

Figure 1. Examined PWR–TES configurations with a dedicated BOP (top) and a dedicated industrial process (bottom). Both configurations employ single-phase to single-phase heat transfer (steam) for TES charging and discharging.

Heat Market and Heat Production Cost Modeling

In this study, we consider a single industrial process or a group of homogeneous heat users (e.g., an industrial park) that have similar heat requirements. Unlike electricity markets, where numerous providers and consumers are interconnected via large-scale grids, heat networks are typically confined to local or regional areas with fewer suppliers and customers [21]. We model a PWR–TES system that maximizes net profit without adhering to a heat-demand schedule. The system can sell any volume of extracted heat at a fixed heat price—up to the 30% dispatch limit—thereby preventing the heat market from serving as an unlimited sink. This setup allows us to focus on which dispatch mode yields the most cost-competitive nuclear heat.

A key economic consideration in determining nuclear heat-production cost is the opportunity cost associated with lost electricity sales [22]. When steam is diverted for cogeneration, electricity output—and the corresponding revenue—may decline. Hence, the heat price must offset both the reduction in electricity revenue and any additional capital or O&M expenses. Formally, the heat-production cost is given by:

$$c_{\text{heat}} = \frac{\Delta C + \Delta\pi_{\text{elec}}}{H_{\text{cogen}}} \tag{1}$$

where c_{heat} is the cost of heat production in \$/MWh$_{(t)}$, ΔC is the increase in NPV of costs due to cogeneration, $\Delta\pi_{\text{elec}}$ is the lost electricity sales revenue due to the cogeneration, and H_{cogen} is the total heat delivered to the process.

The system-NPVs for each scenario are evaluated over a range of heat market prices from \$7 to \$28 /MWh$_{(t)}$, drawing on five years of historical Henry Hub natural gas spot prices (1 January 2020 to 31 December 2024). Converted from \$/MMBtu to \$/MWh$_{(t)}$ the average, minimum, and maximum prices across this period are approximately \$11.67, \$5.08, and \$30.08 /MWh^{t}, respectively. Thus, if the heat market price exceeds the cost of heat production, cogeneration becomes more profitable than the BAU scenario. In this analysis, we determine the heat production cost retrospectively from the dispatch results.

Stochastic Technoeconomic Modeling

In many fields—ranging from meteorology to energy systems planning—there is a need to generate synthetic time series that resemble observed data but extend beyond the single observed realization. Such synthetic series enable scenario testing, uncertainty quantification, and robust decision-making. A key economics derived in heat allocation between electricity generation and process heat is electricity prices from EROCT and PJM. In this study they are represented as synthetic histories through statistical abstraction using autoregressive moving average (ARMA) models. This approach translates historically specific temporal and spatial price patterns into "typical" price histories, preserving the observed average and standard deviation. As a result, this study disentangles the findings from any single market context by emphasizing average electricity prices and volatility in projecting project outcomes (i.e., NPV). Hence, the results remain applicable well beyond the specific markets examined.

The Risk Analysis Virtual Environment (RAVEN) [8] is utilized to generate synthetic histories from the trained reduced order model (ROM). The Holistic Energy Resource Optimization Network (HERON) [9] leverages these synthetic inputs alongside the techno-economic parameters of the energy system to formulate a linear programming or a mixed-integer linear programming (MILP) model. Details regarding the formulation of the ROM algorithm in RAVEN, as well as the capacity and operational optimization in RAVEN-HERON, can be found in Document S1and Document S2, and the open-source model itself [23, 24].

Our economic analysis focuses on the difference in NPV (ΔNPV) between the BAU scenario and the examined scenario. Nuclear capital expenditures (CAPEX) and associated operation and maintenance (O&M) costs are treated as sunk (precommitted). Thus, only TES CAPEX and the resulting changes in electricity and heat sales are included in the analysis. Similarly, the reactor and primary BOP capacities are fixed at 3650 MW$_{(t)}$, while TES capacity is treated as an optimization variable ranging from 1095 to 13,140 MWh$_{(t)}$—equivalent to a 1–12 hour charge duration. For the TES+BOP scenarios, the secondary BOP is also optimized within the same 1095–13,140 MWh$_{(t)}$ range, corresponding to a 1–12 hour discharge duration. Table S2 provides key cost parameters for TES modeling, and Table 1 summarizes the modeled scenarios.

Table 1. Modeled scenarios for different markets assumptions and simulation settings.

Scenario	Market	Heat price [$\$/\mathrm{MWh}_{(t)}$]	Minimum heat supply to industrial customers [$\mathrm{MW}_{(t)}$]	Storage charge duration [hour]
BAU	ERCOT	N/A	N/A	N/A
	PJM	—	—	—
TES + BOP[a]	ERCOT	N/A	N/A	1-12 hours
	PJM	—	—	—
Fixed 30% heat dispatch	ERCOT	7, 10, 13, 16	N/A	1-12 hours
	PJM	—	—	—
Flexible 0-30% heat dispatch	ERCOT	7, 10, 13, 16, 19, 22, 25, 28	N/A	1-12 hours
	PJM	7, 10, 13, 16	—	—
Minimum heat supply[1]	ERCOT	7, 10, 13, 16	500, 800	1-12 hours
	PJM	—	—	—

N/A, not applicable; — denotes the same as above in the table
Both the primary and secondary BOP systems have thermal-to-electricity conversion ratios of 33.64% and 27.53%, respectively. The TES unit is modeled with a 90% round-trip efficiency (RTE) and a periodic boundary condition that fixes the storage level at 75% capacity at both the beginning and end of each simulation.
[a] The heat dispatch modes for both TES+BOP and the minimum heat supply scenarios allow flexible dispatch ranging from 0% to 30%.

RESULTS AND DISCUSSIONS

We present our findings in three parts. First, we compare results from PJM and ERCOT under both fixed-dispatch and flexible dispatch scenarios. Next, we conduct a sensitivity analysis for ERCOT over an extended range of heat prices under the flexible dispatch scenario. Finally, we introduce minimum heat supply requirements into the flexible-dispatch scenario for ERCOT and PJM.

Opportunity Cost of Heat Production

A project's profitability (i.e., NPV > 0) in cogeneration mode hinges on the heat price exceeding the cost of heat production, which must also account for the capital expenditure of the TES. Once the TES size is fixed, the system decides whether to allocate nuclear steam to electricity generation or to heat production by comparing the electricity price to a 'threshold price' ($\mathrm{P}_{\mathrm{THR}}$). When the electricity price is above $\mathrm{P}_{\mathrm{THR}}$, the plant shifts to electricity generation; otherwise, it supplies heat. During optimization, $\mathrm{P}_{\mathrm{THR}}$. indicates the electricity price threshold at which the revenue from selling heat equals that from generating electricity, taking into account relevant BOP efficiency (η_{THR}) and TES round-trip efficiency (η_{RTE}) are factored in. Figure 2 illustrates these relationships, highlighting that higher heat prices tend to raise $\mathrm{P}_{\mathrm{THR}}$ (thus extending cogeneration periods), while improved BOP efficiencies lower $\mathrm{P}_{\mathrm{THR}}$ (favoring electricity generation). Conversely, higher TES round-trip efficiency raises $\mathrm{P}_{\mathrm{THR}}$, making heat dispatch more appealing if the system can recover more value from stored energy.

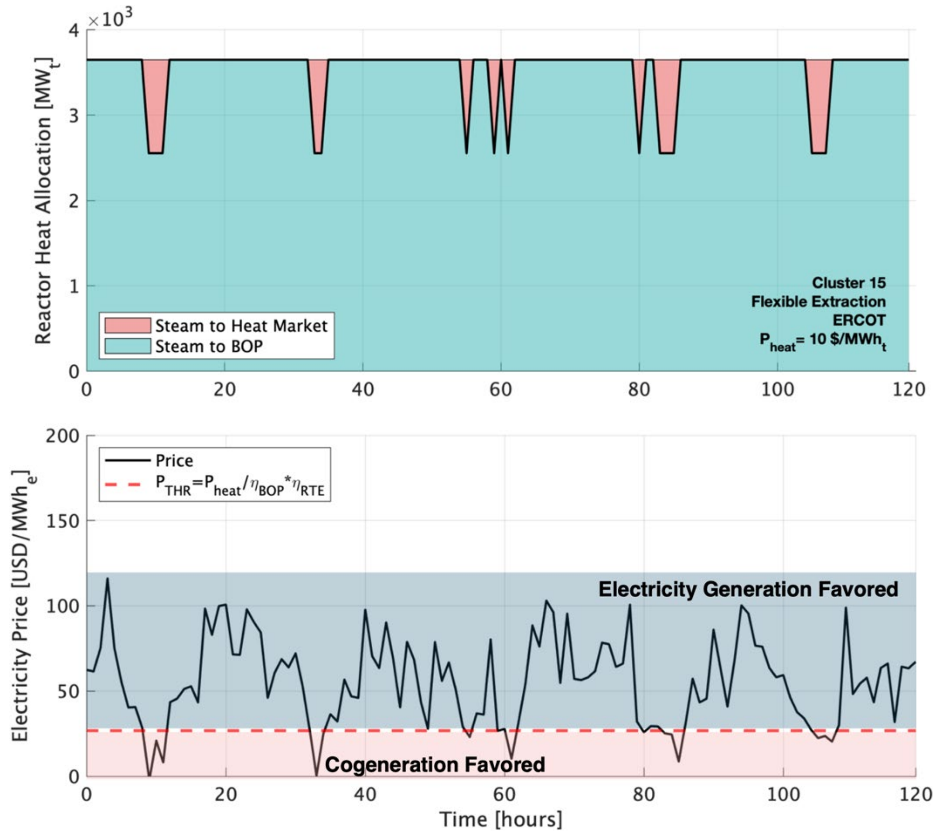

Figure 2. Example of reactor heat dispatch at a \$10/$\mathrm{MWh}_{(t)}$ heat price under the flexible dispatch scenario (top), alongside the corresponding hourly electricity price (bottom).

The threshold price ($\mathrm{P_{THR}}$) indicates when the revenue from selling heat equals that of generating electricity, after accounting for BOP efficiency (η_{BOP}) and TES RTE (η_{RTE}). If the electricity price exceeds $\mathrm{P_{THR}}$, electricity sales are more profitable; if it falls below, heat sales become the economically preferred option.

When cogeneration proceeds without considering P_{THR} (as in the fixed dispatch scenario; see the left panels of Figure S4), the opportunity cost can rise because of forgone electricity sales. This can drive NPV into negative territory if electricity prices are high and heat prices fail to compensate for the lost revenue. However, as illustrated in the right panel of Figure S4, fixed-dispatch cogeneration can still yield a positive NPV when heat prices are high or electricity prices are low, thereby keeping the opportunity cost modest. These findings suggest that industrial applications—where demand is steady and scheduling flexibility is limited—require sufficiently high heat prices to justify cogeneration. Alternatively, improvements in TES round-trip efficiency or in heat-transfer processes (e.g., via the IHX) can help optimize overall performance and profitability.

Table S3 summarizes the capacity combinations that yield the highest NPV in each scenario. In ERCOT, adding TES and a secondary BOP configuration (compared to the BAU baseline) yields the best NPV, largely because it enables electricity-market energy arbitrage. By contrast, in PJM the same configuration results in a negative NPV. This difference arises from: The relatively low thermal-to-electric conversion efficiency of the secondary BOP (27.53%) than the primary BOP; Lower average electricity prices and lower price volatility in PJM, which make it harder to offset the additional capital costs for TES.

Interestingly, in ERCOT, we observe a sudden increase in TES capacity (see Figure S3), largely because TES capital costs are small relative to the total NPV (see Figure S3). As a result, certain ARMA-simulated price scenarios can yield higher energy sales with a larger TES than with a smaller one—even without a direct causal relationship—despite controlling for the mean and standard deviation of electricity prices. To mitigate this effect, incorporating nuclear CAPEX as a balancing factor may help minimize distortions arising from ARMA sampling.

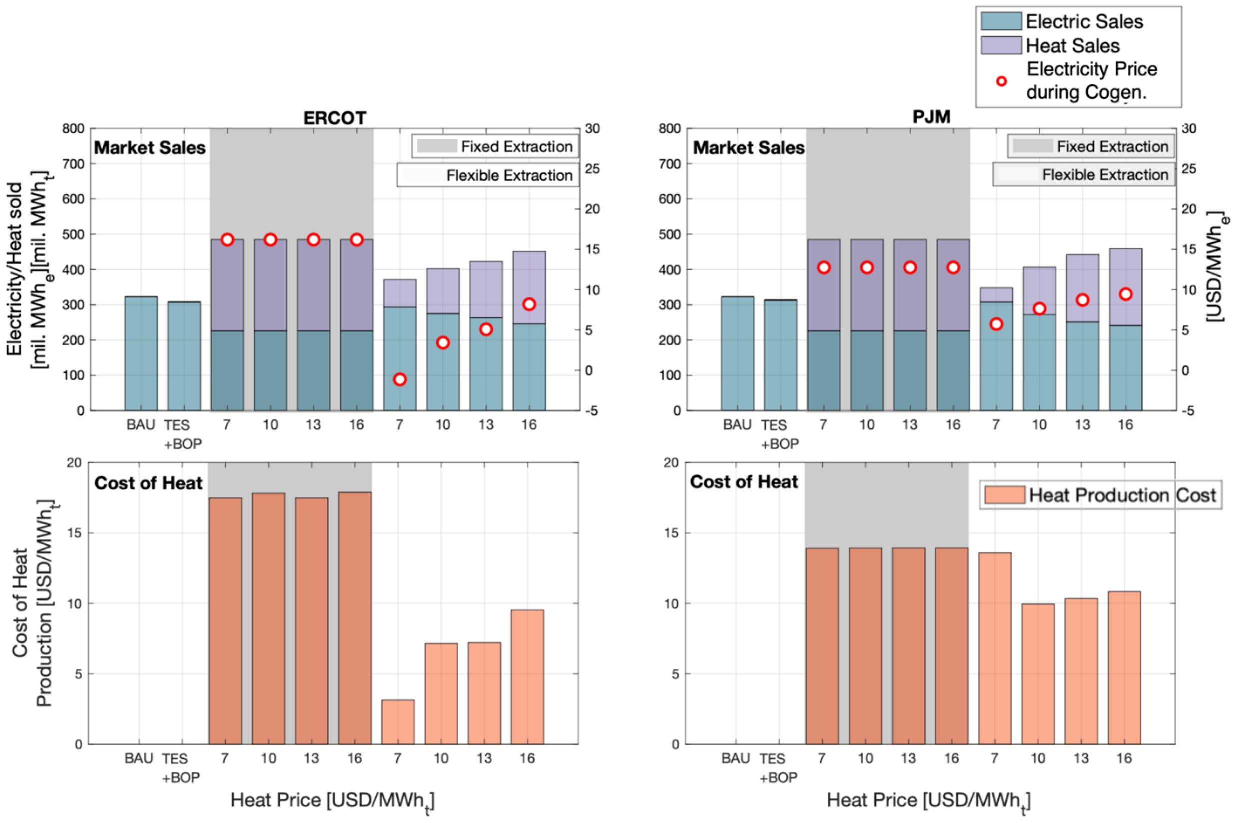

Figure 3. Impact of heat dispatch mode on electricity and heat transactions (top) and heat production cost (bottom) in ERCOT and PJM.
The gray-shaded region at 7–16 $\mathrm{MWh}_{(t)}$ represents the fixed 30% dispatch scenario, while the unshaded 7–16 $\mathrm{MWh}_{(t)}$ range corresponds to the flexible heat dispatch scenario. The "Electricity Price during Cogen." refers to the average electricity price at times when reactor heat is diverted for cogeneration. For the BAU and TES+BOP scenarios, the heat-production graph is empty because these scenarios only generate electricity (i.e., no heat cogeneration). The corresponding optimal sizing and NPV components are provided in Figure S3.

Figure 3 provides a granular view of total electricity and heat transactions, with heat production costs, under both fixed-dispatch (gray shaded area) and flexible-dispatch scenarios at varying heat prices. Thus, one can infer that whenever heat production cost falls below the given heat price, the project's NPV is positive. In the flexible-dispatch scenario, production costs in ERCOT remain below the respective heat price for all examined levels, whereas in PJM, the heat price needs to reach about \$10/MWh$_{(t)}$ for NPV to turn positive. In the top panels of Figure 3, red markers indicate the average electricity price during hours of heat production.

Under a fixed dispatch setup, heat is produced continuously, making the average electricity price during cogeneration essentially the annual average. Because ERCOT's overall electricity prices are higher and more volatile, the average electricity price during fixed dispatch cogeneration is higher in ERCOT than in PJM. Under flexible dispatch, by contrast, the plant schedules heat production during hours of lower electricity prices, causing the average electricity price for cogeneration to drop as low as −\$1.10 to \$8.19/MWh$_{(t)}$ in ERCOT and \$5.70 to \$9.40/MWh$_{(t)}$ in PJM. Consequently, markets with significant price volatility—like ERCOT—offer more profitable opportunities for industrial heat consumers who can adapt their demand. In ERCOT, heat production costs show little variation between \$10 and \$13/MWh$_{(t)}$, remaining around \$8/$MWh_{(t)}$. This is primarily due to an overshooting in TES capacity, driven by ARMA sampling variations. Reducing TES capacity to 1,094 $MWh_{(t)}$ lowers the heat production cost to \$5.1/$MWh_{(t)}$, aligning it more consistently with the heat production cost trend.

When electricity prices are generally stable and stay above the heat price, there is less incentive to switch steam allocation from electricity generation to heat production. This explains why, in PJM at a lower \$7/$MWh_{(t)}$ level, only a small volume of heat is extracted, as the opportunity cost of forgone electricity sales remains high. Another important consideration is that the cost-effectiveness of cogeneration is directly tied to heat prices. Any additional cogeneration beyond the level of heat sales shown in the top panel of Figure 3 in a lower NPV compared to the BAU scenario.

Operational Constraints and Market Signals

The variable (0–30%) heat dispatch mode demonstrates flexibility in adjusting heat production to different heat-price levels, balancing operational constraints and market signals. Our next objective is to determine the heat-price range at which a PWR–TES system can sustain a fixed level of heat dispatch cost-effectively. We investigate this by conducting a sensitivity analysis in ERCOT, spanning heat prices from \$7/$MWh_{(t)}$ to \$28/$MWh_{(t)}$. In cases where industrial processes require a continuous heat supply, identifying this threshold could offer data-informed guidance for aligning the plant's output with an appropriately sized industrial load.

Figure 4 presents the results for optimal electricity and heat production, while the corresponding capacity and NPV details are provided in Figure S5 and Table S4. At \$28/$MWh_{(t)}$, about 27% of the reactor's heat is economically diverted to cogeneration. Although higher heat prices could push dispatch closer to 30%, this is not cost-effective relative to historical natural-gas benchmarks (e.g., a \$30.08/$MWh_{(t)}$ peak at Henry Hub). Consequently, a fixed 30% dispatch rate is suboptimal across most tested heat-price levels (from \$7 to \$16/$MWh_{(t)}$), where the cost-optimal dispatch rises from 9% to 24%. Beyond \$16/$MWh_{(t)}$, dispatch rates rise more gradually, reaching around 27% at \$28/$MWh_{(t)}$. This finding suggests that nuclear plant owners may choose to limit cogeneration to below 30% heat dispatch, thereby maximizing revenues from high electricity-price periods. For instance, in this scenario, an NPP could negotiate a bilateral contract to supply about 24% of its thermal power (on average) at \$16/$MWh_{(t)}$.

An alternative pathway to increase total heat production would involve diverting more than 30% of the reactor's output to non-electric uses. Certain advanced reactor technologies (e.g., the Xe-100) can feasibly direct up to 60% of their capacity to heat applications [25], and recent studies show that coupling with TES can raise this even further—potentially to 100% [26]. However, our analysis focuses on retrofitting existing light-water reactors, where large modifications beyond 30% dispatch would entail major additional costs and engineering challenges, placing such scenarios beyond our scope.

We observed that at a heat price of \$25 /$MWh_{(t)}$, the NPV exceeds that of the BOP+TES scenario (Table S4). Although such a high price may not be guaranteed, the energy arbitrage opportunity in electricity markets is likely to diminish if additional storage providers enter and increase market saturation. Under these conditions, even lower heat-price levels could provide sufficient incentives for the PWR-TES system to serve industrial heat markets, rather than focusing exclusively on electricity production.

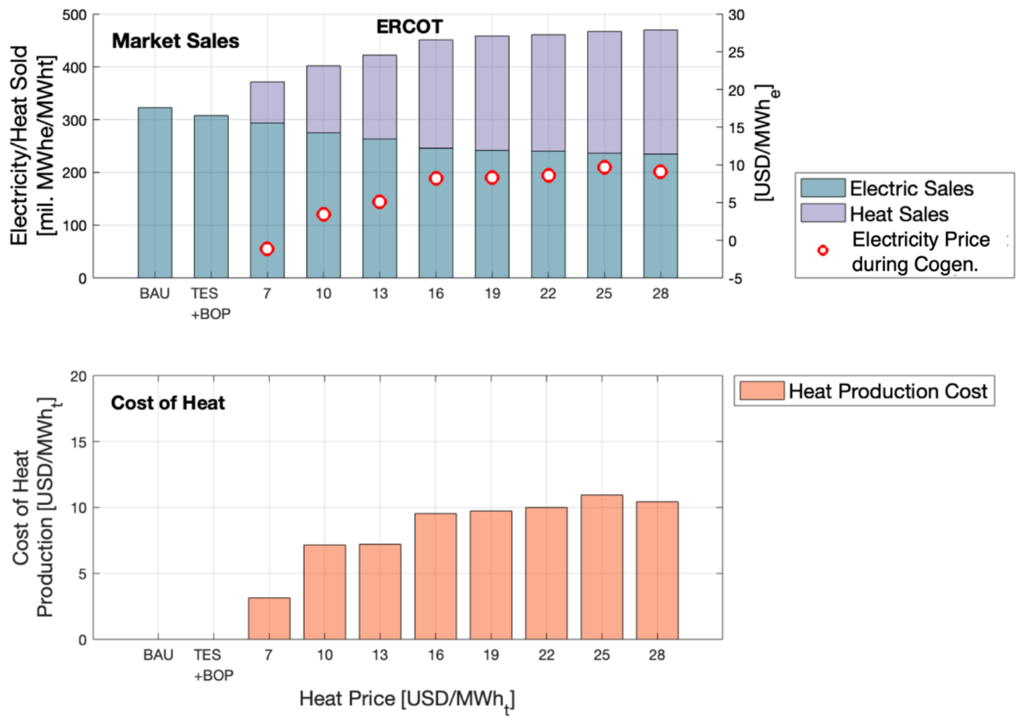

Figure 4. Impact of heat price level on electricity and heat transactions (top) and heat production cost (bottom) in ERCOT.

The 'Electricity Price during Cogen.' refers to the average electricity price at times when reactor heat is diverted for cogeneration. All presented results correspond to the flexible heat-dispatch scenario. For the BAU and TES+BOP scenarios, the heat-production graph is empty because these scenarios only generate electricity (i.e., no heat cogeneration). The corresponding optimal sizing and NPV components are provided in Figure S5.

Minimum Heat Supply Requirement and Heat Production Cost

To explore how minimum heat supply requirements affect PWR–TES systems, we consider a scenario in which the plant must serve an industrial park with aggregated heat demands. Specifically, we focus on 500 $MW_{(t)}$ and 800 $MW_{(t)}$ minimum supply levels—equivalent to 14% and 22% fixed heat dispatch, respectively. These represent the baseload portion of that aggregated demand. Figure 5 presents the cost-effective dispatch of heat and electricity, along with resulting heat production costs, under flexible heat dispatch with these minimum supply requirements. Figure S6 provides detailed capacities and NPV outcomes, while Table S5 compares these results to the fully flexible (no requirement) scenario.

Several key insights emerge: First, at higher minimum supply (800 $MW_{(t)}$), the system in ERCOT can no longer align heat production exclusively with periods of low or negative electricity prices, causing its heat production costs to converge with those in PJM (see bottom panel Figure 5). This represents a significant shift from the heat production costs observed in ERCOT without minimum supply requirements (see bottom panel of Figure 3). As a result, the ERCOT cases show a more pronounced average NPV decline of about 4.8%—compared to 2.7% in PJM (see Table S5). Second, at an 800 $MW_{(t)}$ requirement (about 22% heat dispatch), the plant's dispatch becomes largely insensitive to heat-price variations, reducing the flexibility to optimize between electricity and heat market sales. Lastly, at higher heat-price levels, the penalty from minimum supply requirements becomes less evident, reducing the negative impact on NPV.

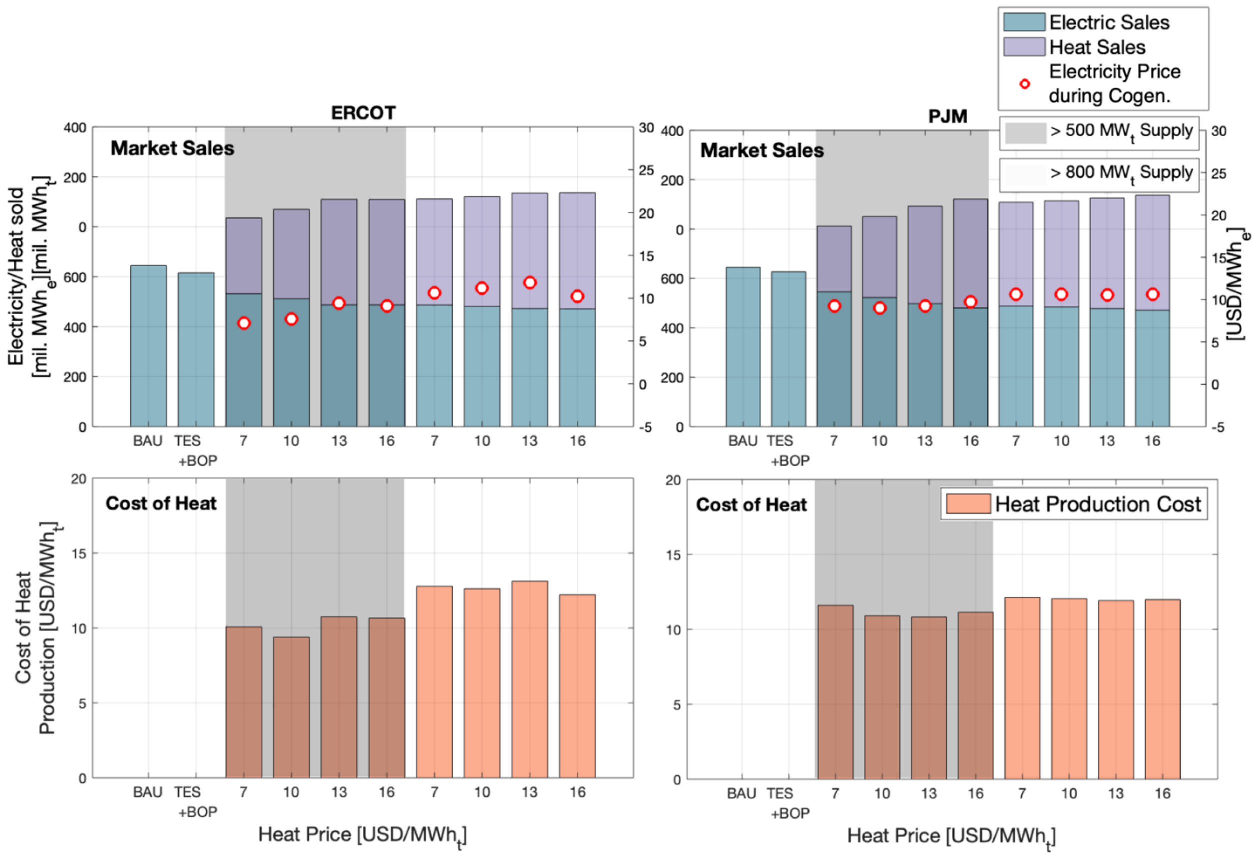

Figure 5. Impacts of minimum heat supply requirements on optimal capacities (top) and NPV components (bottom) under the flexible heat dispatch.
Due to the limited y-axis range, the optimal capacities for the TES+BOP scenario are labeled at the top of the figure. The gray-shaded region at \$7–16/MWh$_{(t)}$ represents the minimum 500 MW$_{(t)}$ heat supply requirement, while the unshaded \$7–16/MWh$_{(t)}$ range corresponds to the represents the minimum 800 MW$_{(t)}$ heat supply requirement. The corresponding optimal sizing and NPV components are provided in Figure S6.

Figure 6 summarizes the above findings in terms of the project's NPV. When industrial users require a minimum heat supply, the PWR-TES plant must charge sufficiently high heat prices to compensate for the revenue forgone from reduced electricity sales during peak-price hours. The trade-offs between heat dispatch modes are less pronounced in PJM, suggesting that the PWR-TES system in PJM has greater flexibility in meeting industrial requirements—due to the relatively low opportunity cost of cogeneration—compared to ERCOT. Moreover, Figure 6 can inform 'go/no-go' decisions. For example, in ERCOT, a 14% minimum dispatch requirement (500 MW$_{(t)}$) causes only a small decrease in NPV (approximately 2%) relative to the fully flexible scenario at a heat price of \$10 MWh$_{(t)}$. If a 2% reduction is acceptable to the supplier, providing heat to the industrial user under these conditions may be both feasible and profitable.

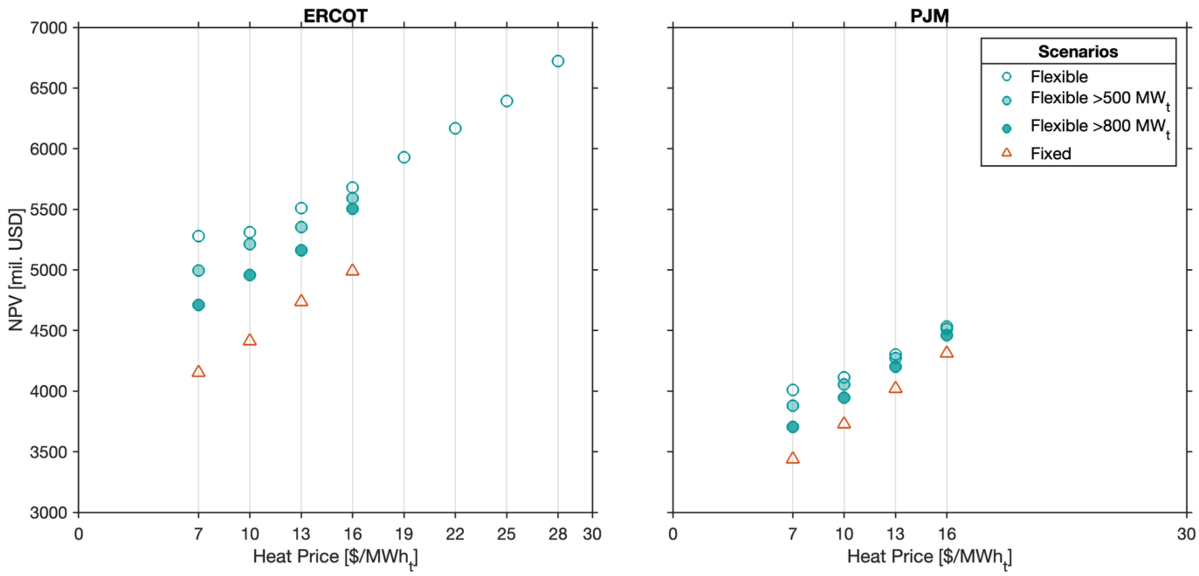

Figure 6. Trade-offs between NPV and minimum heat supply.
As the heat price increases, the resulting NPV values for different minimum heat supply levels (0, 500, and 800 MW$_{(t)}$) illustrate how higher heat prices can help offset potential losses in electricity revenue caused by these operational constraints.

CONCLUSIONS

The viability of nuclear cogeneration depends critically on the relative incentives for both suppliers and consumers, as compared with their available alternatives. For the PWR-TES system, the proposed cogeneration mode must outperform both the BAU baseline and the option of using a TES system solely for electricity market arbitrage. From the consumer's perspective, in turn, the heat production cost from the PWR-TES configuration must be competitive with prevailing natural gas prices (particularly for systems currently relying on gas boilers or combined heat and power units). In this regard, heat dispatch strategies and regional electricity price profiles play pivotal roles in determining the economic viability of nuclear cogeneration.

Results indicate that a fixed heat dispatch mode may not benefit either side. This mode drives the supplier's heat production cost to around $13–17 /MWh$_{(t)}$ —three to five times higher than typical natural gas levels of $3–5 /MWh$_{(t)}$ —primarily because the opportunity cost of diverting steam from electricity sales is substantial. In markets such as ERCOT, characterized by higher electricity prices, such fixed dispatch increases the supplier's foregone electricity revenue. Consequently, industries requiring a continuous heat load under a fixed-dispatch arrangement would either face elevated heat prices or need a larger TES capacity to mitigate mismatches between electricity price spikes and baseload heat demand.

Conversely, a flexible heat dispatch mode—allowing up to 30% of reactor output to be allocated to heat—is found to reduce heat production costs significantly in ERCOT and PJM. Nonetheless, the cost-optimal amount of heat supplied remains below the 30% cap (9–24% of reactor output) at heat-price ranges (e.g., $7–16 /MWh$_t$). Higher baseline heat demands would require additional engineering measures or more flexible reactor technologies.

Finally, in markets with lower electricity prices and reduced volatility (e.g., PJM), the minimum required heat supply has a smaller effect on project NPV. This suggests that operators of existing PWRs in such regions may accommodate industrial heat supply requirements more flexibly and consider cogeneration to create supplementary revenue streams that extend beyond conventional electricity-focused operations.

ACKNOWLEDGMENTS

S.B.C. acknowledges fellowship support from the Idaho National Laboratory (INL). Significant computing resources of the High-Performance Computing Center at the INL were provided, which is supported by the U.S. Department of Energy's Office of Nuclear Energy and the Nuclear Science User Facilities under Contract No. DE-AC07-05ID14517. The authors thank Tyler L. Westover for his advice in the LWRS project that evolved into this research.

AUTHOR CONTRIBUTIONS

S.B.C. performed the modeling, developed the initial experimental design, and contributed to the manuscript production. R.M.S. was lead editor of the manuscript, implemented the storage cost functions to HERON modeling, refined the initial experimental design, and performed analysis of results. S.B.C and R.M.S jointly visualized the results. X.S and T.R.L. advised on analysis of results and reviewed and revised the manuscript.

REFERENCES

[1] EIA. "Natural Gas Weekly Update." https://www.eia.gov/naturalgas/weekly/ (accessed February 1, 2025).

[2] S. H. Kim, T. A. Taiwo, and B. W. Dixon, "The carbon value of nuclear power plant lifetime extensions in the United States," *Nuclear Technology,* vol. 208, no. 5, pp. 775-793, 2022.

[3] N. A. Sepulveda, J. D. Jenkins, F. J. De Sisternes, and R. K. Lester, "The role of firm low-carbon electricity resources in deep decarbonization of power generation," *Joule,* vol. 2, no. 11, pp. 2403-2420, 2018.

[4] S.-B. Cho, A. Epiney, P. Talbot, T. Allen, and X. Sun, "ANALYSIS OF COSTS, BENEFITS, AND IMPLICATIONS OF PATHWAYS IN CAPTURING THE VALUE OF NUCLEAR ENERGY: A CASE STUDY OF TEXAS," in *The Proceedings of the International Conference on Nuclear Engineering (ICONE) 2023.30*, 2023: The Japan Society of Mechanical Engineers, p. 1539.

[5] MIT Energy Initiative, "The Future of Nuclear Energy in a Carbon-Constrained World," 2019. [Online]. Available: https://energy.mit.edu/wp-content/uploads/2018/09/The-Future-of-Nuclear-Energy-in-a-Carbon-Constrained-World.pdf

[6] C. Forsberg, S. Brick, and G. Haratyk, "Coupling heat storage to nuclear reactors for variable electricity output with baseload reactor operation," *The Electricity Journal,* vol. 31, no. 3, pp. 23-31, 2018.

[7] T. Jackowski *et al.*, "Results from the European NC2I-R project on nuclear cogeneration with high temperature reactors-18617," American Nuclear Society-ANS, 555 North Kensington Avenue, La Grange Park …, 2016.

[8] F. Southworth and P. Macdonald, "The next generation nuclear plant (NGNP) project," Idaho National Lab.(INL), Idaho Falls, ID (United States), 2003.

[9] M. Kykal, O. Mazurok, M. Glueckler, C. Herrero Moriana, C. Serin, and V. Korolchuk, "NPP Modification Opportunity by Hydrogen Technologies Implementation: Assessment in Frame of the Euratom Project NPHyCo," in *International Conference on Nuclear Engineering*, 2024, vol. 88216: American Society of Mechanical Engineers, p. V001T01A042.

[10] B. P. Hallbert, "Overview of the US Department of Energy Light Water Reactor Sustainability Program," in *Engineering Assets and Public Infrastructures in the Age of Digitalization: Proceedings of the 13th World Congress on Engineering Asset Management*, 2020: Springer, pp. 932-936.

[11] S. M. Bragg-Sitton *et al.*, "Integrated energy systems: 2020 roadmap," Idaho National Lab.(INL), Idaho Falls, ID (United States); Oak Ridge …, 2020.

[12] U.S. Department of Energy. "Nine Mile Point Begins Clean Hydrogen Production." https://www.energy.gov/ne/articles/nine-mile-point-begins-clean-hydrogen-production (accessed February 1, 2025).

[13] K. Frick, D. Wendt, P. Talbot, C. Rabiti, and R. Boardman, "Technoeconomic assessment of hydrogen cogeneration via high temperature steam electrolysis with a light-water reactor," *Applied Energy,* vol. 306, p. 118044, 2022.

[14] S. J. Root *et al.*, "Comparison of Energy Storage and Arbitrage Options for Nuclear Power," 2023.

[15] L. T. Knighton *et al.*, "Energy Arbitrage: Comparison of Options for use with LWR Nuclear Power Plants," Idaho National Lab.(INL), Idaho Falls, ID (United States), 2021.

[16] D. J. McDowell *et al.*, "A Technical and Economic Assessment of LWR Flexible Operation for Generation and Demand Balancing to Optimize Plant Revenue," Idaho National Lab.(INL), Idaho Falls, ID (United States), 2021.

[17] R. M. Saeed *et al.*, "Multilevel Analysis, Design, and Modeling of Coupling Advanced Nuclear Reactors and Thermal Energy Storage in an Integrated Energy System," Idaho National Lab.(INL), Idaho Falls, ID (United States), 2022.

[18] R. M. Saeed, K. L. Frick, A. Shigrekar, D. Mikkelson, and S. Bragg-Sitton, "Mapping thermal energy storage technologies with advanced nuclear reactors," *Energy Conversion and Management,* vol. 267, p. 115872, 2022.

[19] D. M. Mikkelson, K. L. Frick, C. Rabiti, and S. M. Bragg-Sitton, "Thermal Energy Storage Model Development within the Integrated Energy Systems HYBRID Repository," Idaho National Lab.(INL), Idaho Falls, ID (United States), 2021.

[20] S. Hancock and T. Westover, "Simulation of 15% and 50% Thermal Power Dispatch to an Industrial Facility Using a Flexible Generic Full-Scope Pressurized Water Reactor Plant Simulator," *Energies,* vol. 15, no. 3, p. 1151, 2022.
[21] NEA, "Beyond Electricity: The Economics of Nuclear Cogeneration," OECD Publishing, Paris, France, 2022. [Online]. Available: https://www.oecd-nea.org/upload/docs/application/pdf/2022-07/7363_cogen.pdf
[22] H. Safa, "Economics of District Heating using Light Water Reactors," presented at the NEA/IAEA Workshop on Nuclear Cogeneration, Paris, France, 2013.
[23] Idaho National Laboratory. "RAVEN Github Repository." https://github.com/idaholab/raven (accessed January 22, 2024).
[24] Idaho National Laboratory. "HERON Github Repository." https://github.com/idaholab/heron (accessed January 22, 2024).
[25] E. Mulder and W. Boyes, "Neutronics characteristics of a 165 MWth Xe-100 reactor," *Nuclear Engineering and Design,* vol. 357, p. 110415, 2020.
[26] R. M. Saeed, V. Novotný, S.-B. Cho, A. Shigrekar, C. Otani, and D. Mikkelson, "Design and optimization of flexible decoupled high-temperature gas-cooled reactor plants with thermal energy storage," *Energy Conversion and Management,* 2024.

SUPPLEMENTAL INFORMATION

Expanding Market Opportunities: Cogeneration Strategies for Integrated PWR and Thermal Energy Storage Systems

So-Bin Cho, Rami M. Saeed, Todd Allen, and Xiaodong Sun

Document S1. Statistical Representation of Electricity Prices

In this study, synthetic data is employed to replicate electricity prices in ERCOT and PJM. The synthetic data creation proceeds by: (1) segmenting the year-long data into smaller time intervals, (2) detrending seasonal components via Fourier series, (3) fitting an ARMA model to the residuals within short segments, (4) clustering the ARMA parameters into a reduced set using K-means, and (5) reconstructing synthetic histories that preserve both the deterministic seasonality and the stochastic noise characteristics. The central assumption behind this method is that detrending via a suitable Fourier series and segmenting the data enables stationarity within each short interval. While stationarity may not be perfectly satisfied in practice, the segmentation step typically minimizes non-stationary artifacts. Figure S1illustrates the first four steps of this process.

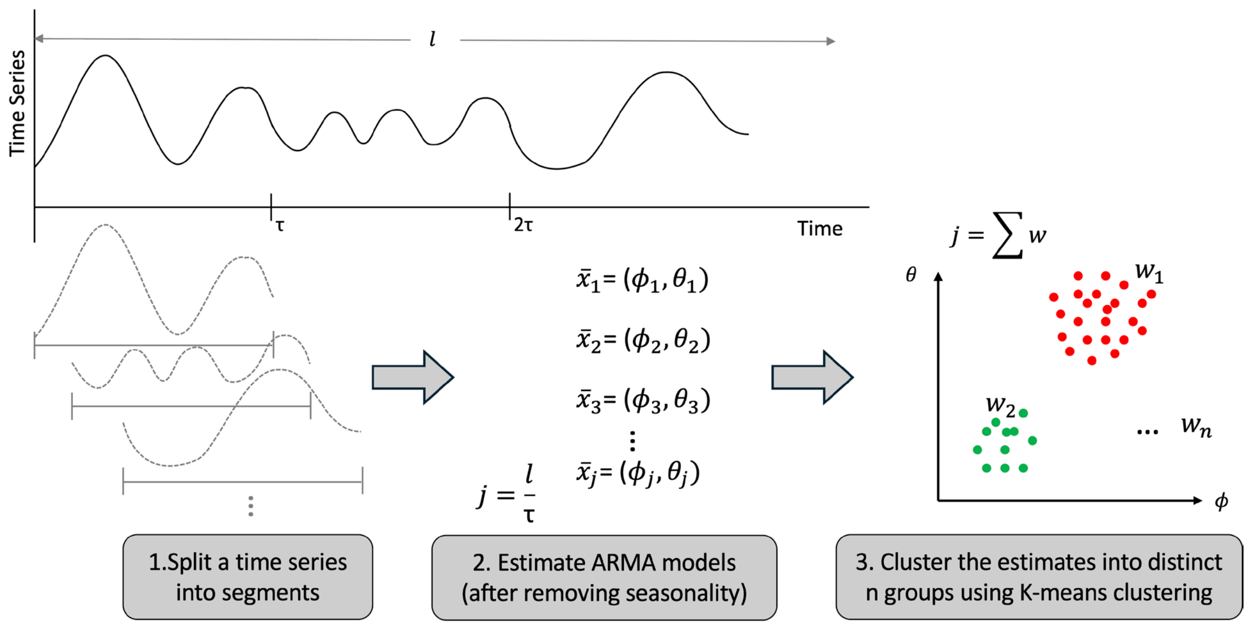

Figure S1. Flowchart illustrating the process of synthetic data training: segmentation, ARMA fitting, and Clustering from left to right.

ARMA models inherently assume stationarity, which requires stationarity assumptions (i.e., constant mean, covariance, and covariance). However Real-world data frequently contain time-varying patterns that may violate stationarity. Hence, a preliminary step is to divide the data into smaller segments, each of which can be treated (approximately) as stationary. Let l is the total length of the time series (e.g., a year's worth of data). The choice of temporal extent **(**segment length**)**, τ is critical. A long τ risks including non-stationary behavior within a single segment, while a short τ might fail to capture the relevant stochastic features within each segment. Dividing the year (l) into segments with temporal extent τ yields j segments.

After removing the fitted periodic signal (i.e., via Fourier detrending), the residuals are assumed to be stationary and used for ARMA modeling. If each segment has its own ARMA parameters, the total number of distinct parameter sets can become quite large. To reduce complexity and adopt a 'reduced-orde" approach, these parameter vectors can be clustered using the standard K-means algorithm. Upon convergence, N representative ARMA parameter sets are obtained. Each cluster is associated with a weight w_i indicating how many segments fall into that cluster, and these weights sum to the total number of segments j, thus preserving the original time coverage for each cluster type. Table S1 provides detailed information used in the synthetic data creation.

Table S1. Synthetic data training and realization settings used in HERON simulations.

Parameter	Description	Value/Range
Project life	Total duration of simulation	30 years
Temporal extent	A dispatch time frame.	120-hour
Temporal resolution	Simulation step size for dispatch	1 hour
ARMA samples	The number of ARMA samples to capture variability	20
Measurement	Input for synthetic history training	ERCOT 2018–2021 LMP data (West LZ) [a] PJM 2018–2021 LMP data (Brandon Shores) [a]
ARMA order	The (p, q) order used for the ARMA model .	p = 1, q = 1 (electricity price) [b]
Removed Fourier seasonality (electricity price)	Periods detrended from the time series before ARMA fitting	24, 168, 199, 350, 365, 674, 730, 1095, 1251, 1752, 2190, 2920, 4380, 8760
Number of clusters	Predefined numbers for *k*-means clustering.	20[c]

a. 2018–2021 data, trained to produce a 30-year horizon (in 4-year repeating blocks)

b. Determined via AIC/BIC model selection.

c. All synthetic inputs must be generated with the same number of clusters for HERON

Document S2. Capacity and Operation Optimization

In examining the optimal capacity and operation of energy systems, the RAVEN workflows generated by HERON involve a bi-level optimization process. A schematic representation of the framework is shown in Figure S2. The outer loop (capacity optimization) aims to determine profit-maximizing capacities by sampling the capacities of all system components. Capacity optimization is carried out using either the optimization mode or the sweep mode. The optimization mode explores the capacity space using an optimizer, such as gradient descent algorithms. On the other hand, the sweep mode conducts comprehensive parametric studies by 'sweeping' across all possible combinations of capacities. As more components are optimized, the sweep mode often requires numerous iterations due to the large dimensionality of the combined capacity space.

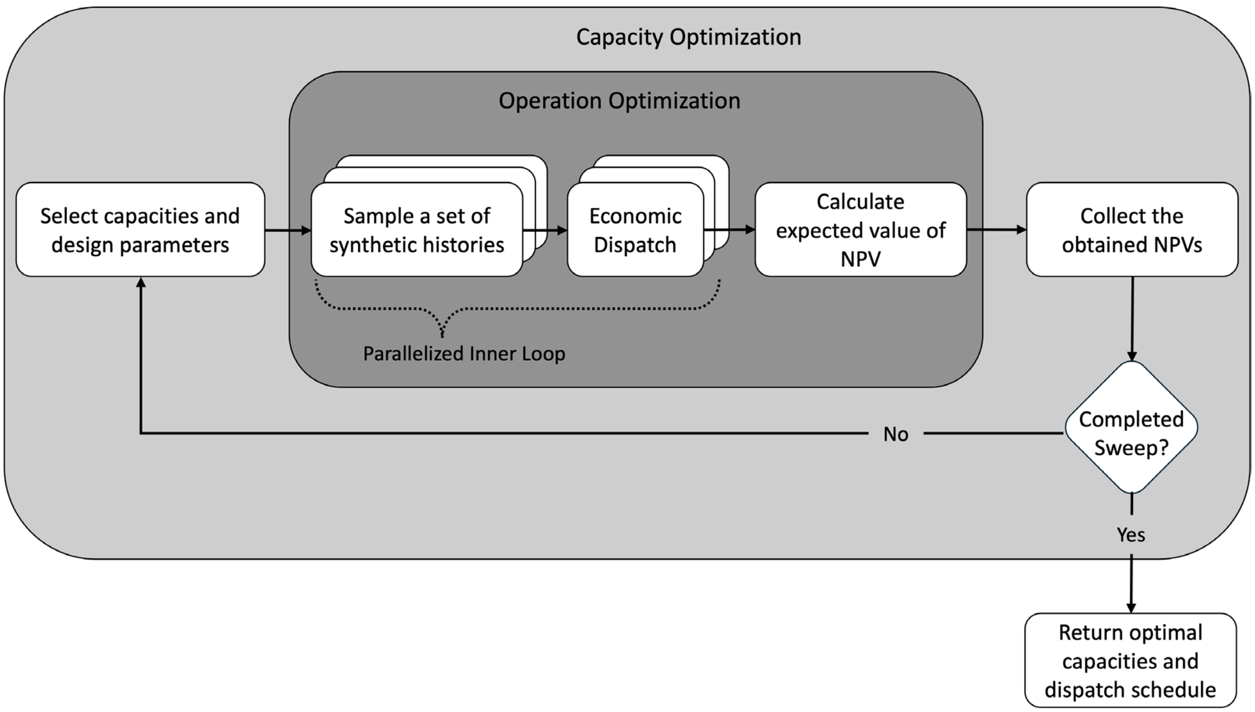

Figure S2. Flowchart for capacity and operation optimization.

The inner loop optimizes the operation of each component for each time step, with the capacity variables from the outer loop held constant. It samples a predefined number of synthetic histories (see ARMA samples in Table S1) to account for uncertainty in temporal profiles. A single sample of a ROM is designed to account for variability in historical data by maintaining identical mean and standard deviation. However, optimal operation can still be influenced by chronological or sequential differences. By optimizing over multiple realizations, the stochastic nature of the optimization problem is included. The primary benefit of RAVEN workflows is their ability to deliver metrics of interest, such as NPV and dispatch-related variables, in terms of expected values.

The objective function of the HERON input in this thesis is to maximize the NPV of net profit by considering both cost and revenue components. This can be expressed as:

$$\text{maximize}\, \pi - c, \qquad \textbf{Equation 1}$$

where π represents revenues from either the electricity or heat market, and c includes both investment costs and operating costs for all components. Since RAVEN generates a set of synthetic histories (through sampling) from a reduced order model, optimization is performed to maximize the expected value of the system. The mathematical expression for the optimization is as follows:

$$M_{opt} = \max_{\vec{C},\vec{D}} \left(\mathbb{E}_{\Omega} \left(M(\vec{C}, \vec{D}, \Omega) \right) \right), \qquad \textbf{Equation 2}$$

where M is the economic metric of interest, specifically net profit in this thesis; M_{opt} is the optimal value of the metric; $\vec{C}$ is the vector containing all component capacities, such as nuclear reactors, storage systems, and BOP; $\vec{D}$ is the vector containing all subcomponent dispatches, such as the generated electricity from BOP for each time step of the program.; Ω represents a single set of evaluations of the synthetic time histories.; $\mathbb{E}_{\Omega}$ is the expected value operator over realizations Ω. For example, if 20 electricity price samples are generated from an ARMA reduced order model, the expected values of net profit will be calculated using E_{Ω}. The input spaces $\vec{C}$ and $\vec{D}$ that provide M_{opt} are the optimal capacity and operation, respectively. Equation 2 is reformulated using the net profit definition (Equation 1):

$$M_{opt} = \max_{\vec{C},\vec{D}} \left(\mathbb{E}_{\Omega} \left(\pi(\vec{C}, \vec{D}, \Omega) - c(\vec{C}, \vec{D}, \Omega) \right) \right) \qquad \textbf{Equation 3}$$

The bi-level optimization problem reformulated from Equation 3 is provided in Equation 4.

$$M_{opt} = \max_{\vec{C}} \left(\mathbb{E}_{\Omega} \left(\max_{\vec{D}} \left(\pi(\vec{C}, \vec{D}, \Omega) - c(\vec{C}, \vec{D}, \Omega) \right) \right) \right) \qquad \textbf{Equation 4}$$

Recognizing that the operation optimization is conducted only over a set of shorter segments (e.g., synthetic electricity prices trained with 24-hour or 120-hour temporal extent), rather than spanning the entire 8760 hours, Equation 4 can be rearranged so that each cluster n is weighted by the number of segments it represents as follows:

$$M_{\text{opt}} = \max_{\vec{C}} \left\{ \underbrace{-c_{\text{cap}}(\vec{C})}_{\text{Capital cost (one-time)}} + \sum_{y=1}^{Y} \left[\underbrace{-c_{\text{fOM}}(\vec{C})}_{\text{Fixed O\&M cost (annual)}} + \mathbb{E}_{\Omega} \left[\max_{\vec{D}} \left\{ \underbrace{\sum_{n=1}^{N} w_n \left(\pi_n(\vec{C}, \vec{D}, \Omega) - c_{\text{vOM},n}(\vec{C}, \vec{D}, \Omega) \right)}_{\text{Weighted revenue minus variable O\&M}} \right\} \right] \right] \right\} \qquad \textbf{Equation 5}$$

The capital cost of a component is not annualized and is paid as a lump sum at project time y = 0, while the fixed O&M and the annual revenues continue to occur each year throughout the project life, up to year *Y*. Here, *N* is the total number of clusters, and w_n scales the cluster-level net profit to reflect the true fraction of the year represented in cluster *n*; the expressions c_{cap} and c_{fOM} correspond to the total capital cost and annual fixed O&M, respectively, as functions of the capacity vector $\vec{C}$; The total variable O&M cost in cluster *n*, denoted as $c_{\text{vOM},n}$, depends on $\vec{C}$, $\vec{D}$, and Ω. The main idea is to separate terms that are not associated with the temporal extent of each cluster (e.g., CAPEX, fixed O&M) from those that scale with the cluster weight w_n (e.g., energy-market revenues). In other words, CAPEX and annual fixed O&M are not "partial". Therefore, these costs appear outside the $\sum_{n=1}^{N} w_n[\ldots]$ portion.

In this system, we consider a TES setup consisting of a charger (intermediate heat exchanger and piping), and a storage tank. In addition, for the TES+BOP scenario, the system also includes a discharger (intermediate heat exchanger and piping, plus the secondary balance of plant system for the BOP+TES scenario). We assume that the fixed and variable O&M)costs for the TES system are negligible for existing nuclear power plant (NPP) projects. Revenue is derived from both heat and electricity sales. Under these assumptions, Equation 5 simplifies to Equation 6.

$$M_{\text{opt}} = \max_{\vec{C}} \left\{ \underbrace{-c_{\text{cap}}(\vec{C})}_{\text{Capital cost (TES)}} + \sum_{y=1}^{Y} \left[\mathbb{E}_{\Omega} \left[\max_{\vec{D}} \left\{ \underbrace{\sum_{n=1}^{N} w_n \left(P^{\text{elec}}(\Omega) E(\vec{C}, \vec{D}, \Omega) + P^{\text{heat}}(\Omega) H_{cogen}(\vec{C}, \vec{D}, \Omega) \right)}_{\text{Weighted electricity \& heat sales revenue}} \right\} \right] \right] \right\} \quad \textbf{Equation 6}$$

where P^{elec} and P^{heat} are the prices for electricity and heat, respectively. E And H_{cogen} represent the electricity and heat transactions, respectively, in their corresponding markets. Note that the fixed costs for the TES system are incorporated using cost functions that reflect economies of scale. The cost c_{cap} of each subsystem capacity C is modeled as:

$$c_{\text{cap}}(C_{TES\ subsystem}) = A \left(\frac{D}{D'} \right)^{x} \quad \textbf{Equation 7}$$

where A is a reference cost, D' is a reference capacity, D is the capacity under consideration, and x is the scaling factor. Table S2 presents the parameters of A, D', and x for each TES subsystem [17].

Table S2. Cost function parameters for the three superset models for the PWR-TES coupling.

Superset model	*A*	*D'*	*x*
Charger	2,964,480.3	72.57 $MW_{(t)}$	0.95986969
TES	36,452,122.8	435.42 $MWh_{(t)}$	0.83976343
Discharger*	10,896,427.1	72.57 $MW_{(t)}$	0.69183838

* In the cost function for the discharge superset that we used, turbine costs for electricity generation are included. Thus, when stored heat is sold directly to industrial customers without a secondary BOP, no capital expenditure for the delivering heat to customers via the discharge superset were assigned. This setting implicitly assumes that PWR-TES systems are integrated into an existing heat delivery network. Depending on the ownership of the PWR-TES systems and the structure of the heat market, detailed cost functions for the discharge superset (without a BOP consideration) need to be evaluated.

Table S3. HERON optimized cases for the PWR-TES coupling.

Market	Scenario	Heat price [$/MWh_{(t)}$]	TES, 2nd. BOP ($MWh_{(t)}$, $MW_{(t)}$)		Mean NPV [mil. USD]	Delta NPV[1] [mil. USD], [% difference]	
ERCOT	BAU	N/A	0	0	5172	N/A	N/A
	TES + BOP	-	13140	12045	6326	1154	22.3
	Fixed 30% heat dispatch	7	1095	0	4153	-1019	-19.7
		10	1095	0	4413	-760	-14.7
		13	1095	0	4736	-436	-8.4
		16	3285	0	4989	-184	-3.6
	Flexible 0-30% heat dispatch	7	1095	0	5278	106	2.0
		10	2190	0	5309	137	2.6
		13	1095	0	5510	337	6.5
		16	1095	0	5679	506	9.8
PJM	BAU	N/A	0	0	4111	N/A	N/A
	TES + BOP	N/A	4380	4380	4103	-7	-0.2
	Fixed 30% heat dispatch	7	1095	0	3440	-671	-16.3
		10	1095	0	3729	-381	-9.3
		13	1095	0	4020	-91	-2.2
		16	1095	0	4311	201	4.9
	Flexible 0-30% heat dispatch	7	1095	0	4011	-100	-2.4
		10	1095	0	4113	2	0.1
		13	1095	0	4301	190	4.6
		16	1095	0	4533	423	10.3

2nd. BOP, secondary BOP.
The capacities across all cases for the PWR, primary BOP, and charger superset are fixed at 3650 $MW_{(t)}$, 3650 $MW_{(t)}$, and 1095 $MW_{(t)}$, respectively.

[1] *Delta NPV: $\Delta NPV = NPV_{case} - NPV_{BAU}$*

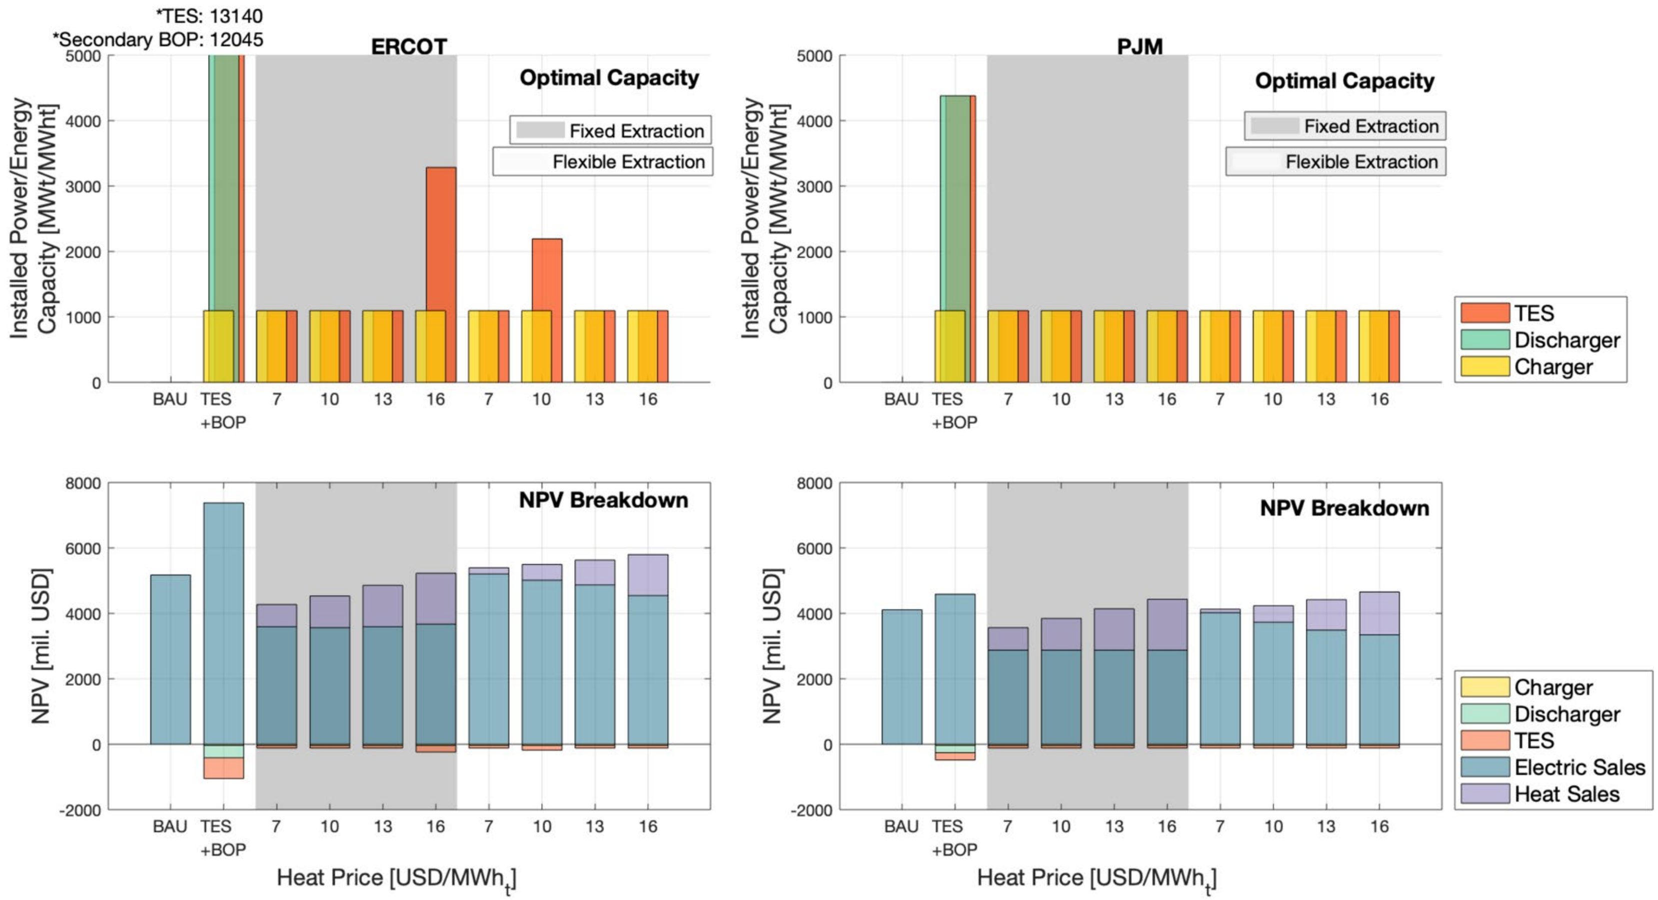

Figure S3. Impacts of heat price level on optimal capacities (top) and on NPV components (bottom).
The gray-shaded region at \$7–16/MWh$_{(t)}$ represents the fixed 30% dispatch scenario, while the unshaded \$7–16/MWh$_{(t)}$ range corresponds to the flexible heat dispatch scenario.

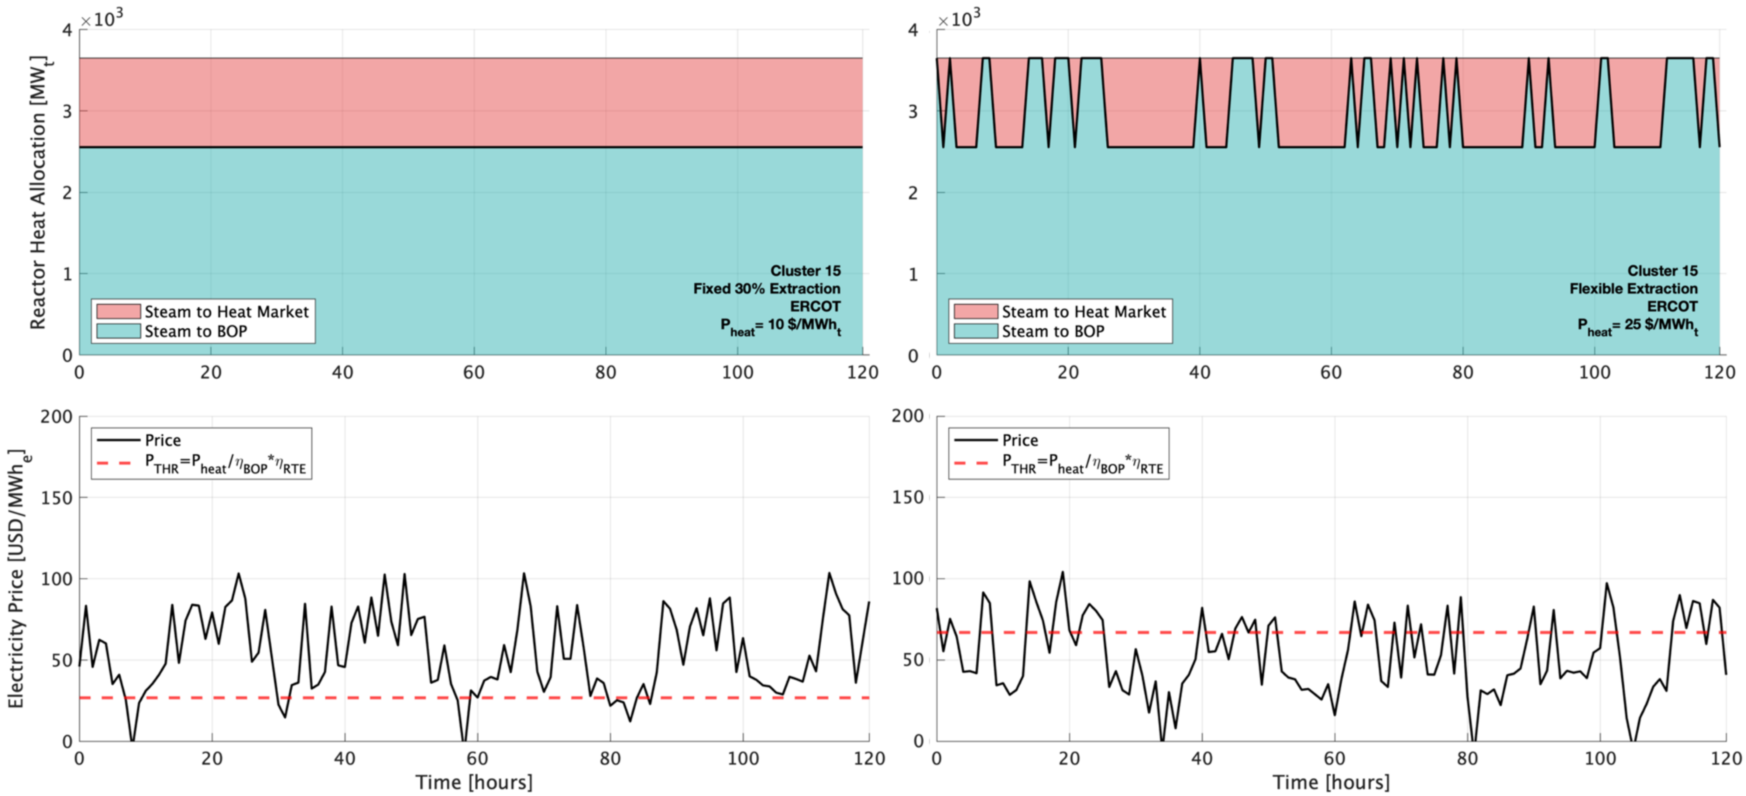

Figure S4. Example of reactor heat dispatch at $10/MWh (left) and $25/MWh (right) heat price, along with the corresponding hourly electricity price.

The left panels show the fixed dispatch scenario, and the right panels show the flexible scenario. The bottom panel displays the corresponding hourly electricity price. The threshold price (P_{THR}) indicates when the revenue from selling heat equals that of generating electricity, after accounting for BOP efficiency (η_{BOP}) and TES RTE (η_{RTE}). If the electricity price exceeds P_{THR}, electricity sales are more profitable; if it falls below, heat sales become the economically preferred option.

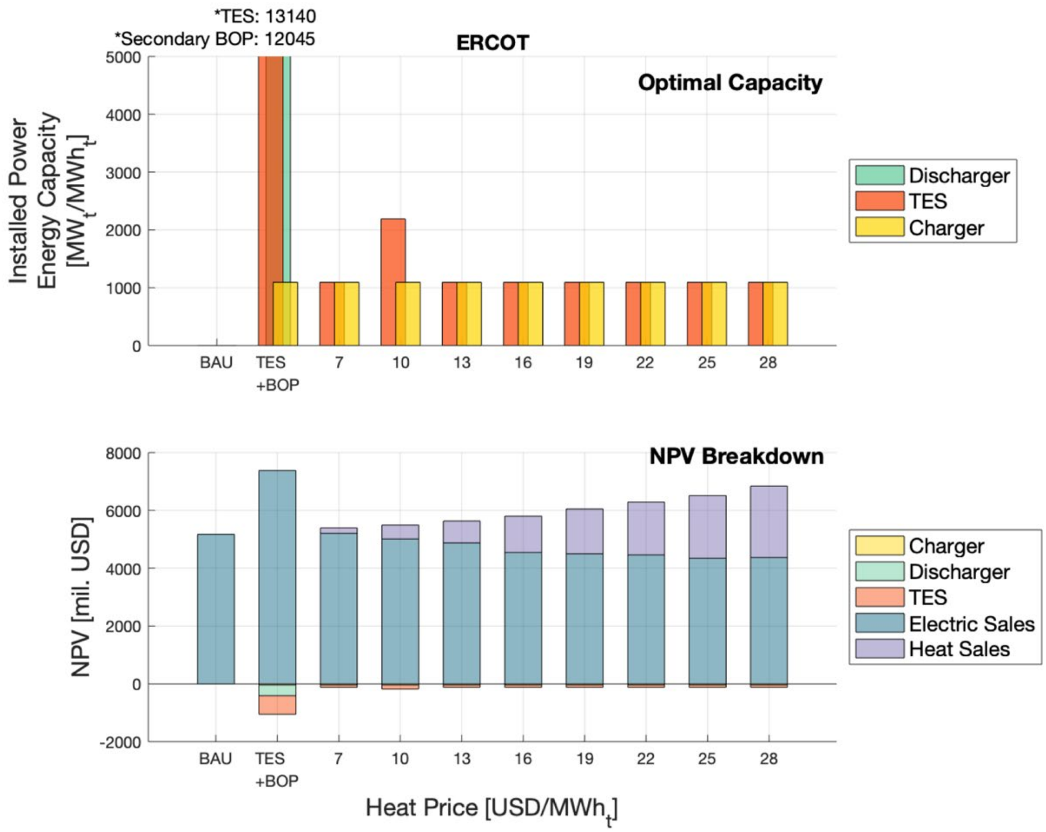

Figure S5. Impacts of the heat price level on optimal capacities (top) and on NPV components (bottom) in ERCOT.

All presented results correspond to the flexible heat-dispatch scenario. For the BAU and TES+BOP scenarios, the heat-production graph is empty because these scenarios only generate electricity (i.e., no heat cogeneration). Due to the limited y-axis range, the optimal capacities for the TES+BOP scenario are labeled at the top of the figure.

Table S4. Impact of heat price levels on optimal TES sizing and NPV in ERCOT.

Market	**Scenario**	**Heat price** [$/MWh$_{(t)}$]	**TES, 2nd. BOP** (MWh$_{(t)}$, MW$_{(t)}$)		**Mean NPV** [mil. USD]	**Delta NPV[1]** [mil. USD], [% difference]	
ERCOT	BAU	-	-	-	5172	-	-
	TES + BOP	-	13140	12045	6326	1154	22.3
	Flexible 0-30% heat dispatch	7	1095	-	5278	106	2.0
		10	2190	-	5309	137	2.6
		13	1095	-	5510	337	6.5
		16	1095	-	5679	506	9.8
		19	1095	-	5930	757	14.6
		22	1095	-	6168	996	19.2
		25	1095	-	6392	1220	23.6
		28	1095	-	6722	1549	30.0

2nd. BOP, secondary BOP.
The capacities across all cases for the PWR, primary BOP, and charger superset are fixed at 3650 MW$_{(t)}$, 3650 MW$_{(t)}$, and 1095 MW$_{(t)}$, respectively.

[1] *Delta NPV: $\Delta NPV = NPV_{case} – NPV_{BAU}$*

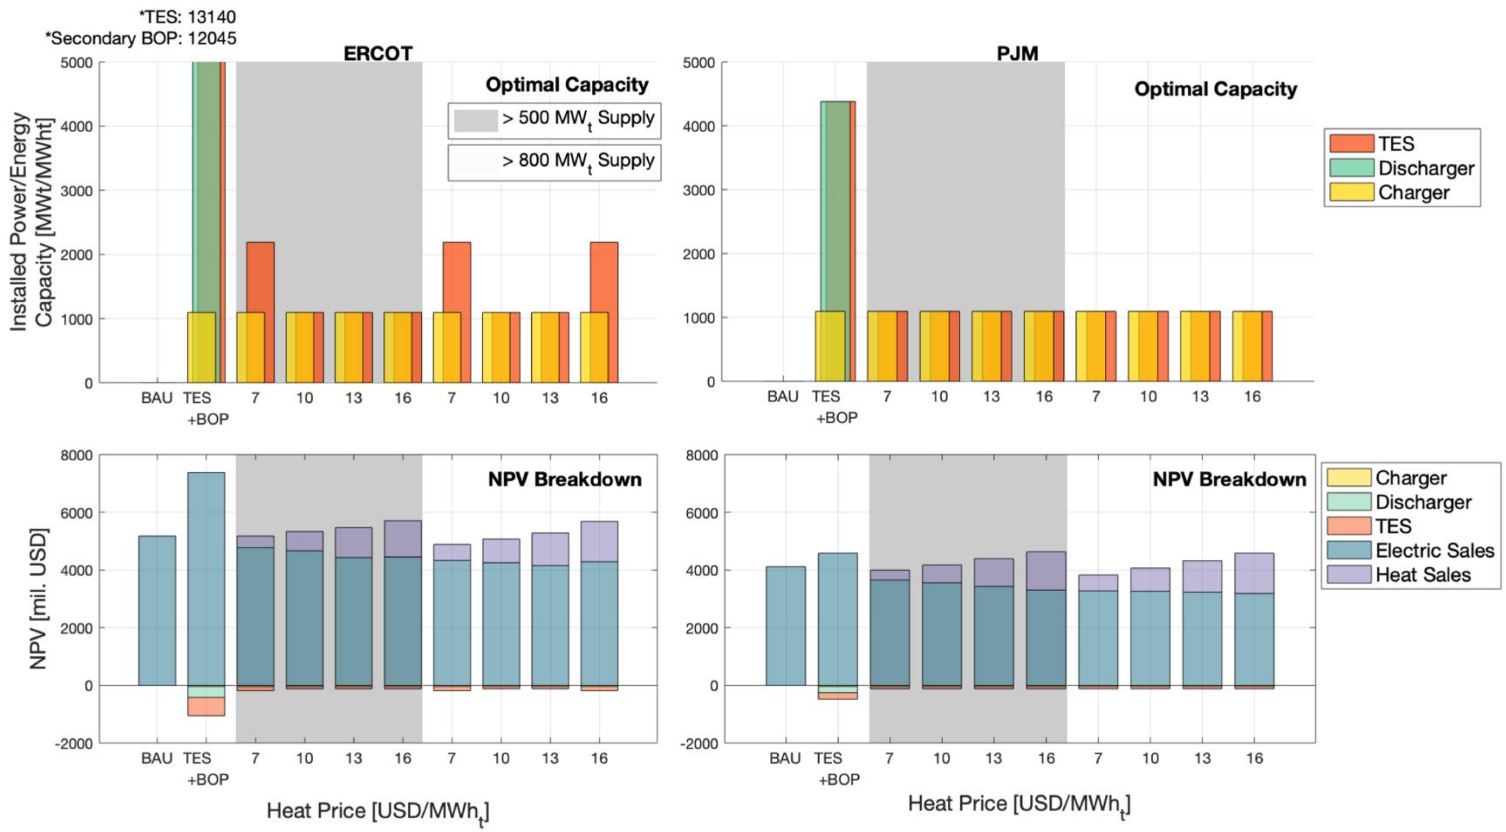

Figure S6. Impacts of minimum heat supply requirements on optimal capacities (top) and NPV components (bottom) under the flexible heat dispatch.

Due to the limited y-axis range, the optimal capacities for the TES+BOP scenario are labeled at the top of the figure. The gray-shaded region at \$7–16/MWh$_{(t)}$ represents the minimum 500 MW$_{(t)}$ heat supply requirement, while the unshaded \$7–16/MWh$_{(t)}$ range corresponds to the represents the minimum 800 MW$_{(t)}$ heat supply requirement.

Table S5. Impact of Minimum Heat Supply Levels on Optimal Sizing and NPV in ERCOT and PJM

Market	Scenario	min. heat supply [$MW_{(t)}$]	Heat price [\$/$MWh_{(t)}$]	TES [$MWh_{(t)}$]	Mean NPV [mil. USD]	Delta NPV[1] [mil. USD], [% difference]	
ERCOT	Flexible 0-30% heat dispatch	0	7	1095	5278	106	2.0
			10	2190	5309	137	2.6
			13	1095	5510	337	6.5
			16	1095	5679	506	9.8
	Minimum heat supply	500	7	2190	4995	-178	-3.4
			10	1095	5213	40	0.8
			13	1095	5354	181	3.5
			16	1095	5592	420	8.1
		800	7	2190	4710	-463	-8.9
			10	1095	4957	-215	-4.2
			13	1095	5162	-11	-0.2
			16	2190	5504	331	6.4
PJM	Flexible 0-30% heat dispatch	0	7	1095	4011	-100	-2.4
			10	1095	4113	2	0.1
			13	1095	4301	190	4.6
			16	1095	4533	423	10.3
	Minimum heat supply	500	7	1095	3879	-231	-5.6
			10	1095	4055	-56	-1.4
			13	1095	4272	161	3.9
			16	1095	4513	402	9.8
		800	7	1095	3706	-405	-9.9
			10	1095	3946	-165	-4.0
			13	1095	4201	90	2.2
			16	1095	4462	351	8.5

min., minimum.

The capacities across all cases for the PWR, primary BOP, and charger superset are fixed at 3650 $MW_{(t)}$, 3650 $MW_{(t)}$, and 1095 $MW_{(t)}$, respectively.

[1] *Delta NPV: $\Delta NPV = NPV_{case} - NPV_{BAU}$*