

Benders Decomposition with Partial Non-Anticipativity Relaxation for Multi-Stage Stochastic Clean Energy Transition Planning

Ahmet Emir Şener^a, Burak Kocuk^{b,*}, Tuğçe Yüksel^b

^aDepartment of Industrial Engineering, Boğaziçi University, Istanbul, Turkey

^bFaculty of Engineering and Natural Sciences, Sabancı University, Istanbul, Turkey

Abstract

We study clean energy transition planning for campus-scale integrated electricity-heat systems under both strategic level and operational level uncertainties. We formulate a multi-stage stochastic mixed-integer program that jointly optimizes investment and operational decisions for renewable generation, storage, and heat-transfer technologies whose costs and efficiencies evolve stochastically across stages. To account for short-term operational uncertainty, we further derive a robust reformulation based on box uncertainty sets for demand, renewable generation, and heat-transfer performance. To solve the resulting large-scale model, we develop a Benders decomposition algorithm with partial non-anticipativity relaxation. Investment variables and their non-anticipativity constraints are retained in the master problem while operational variables are assigned to scenario-wise subproblems where their non-anticipativity constraints are relaxed. Non-anticipativity of operational variables is restored only at termination through a smaller linear program. We prove that this correction step can only increase the objective function value by a finite bound. We improve the computational performance of the algorithm developed with valid inequalities and a two-phase cut-addition strategy. Using the proposed approach, we solve the Middle East Technical University campus case study at high temporal resolution within a reasonable computational budget, which is not otherwise possible with the extensive form or the classical Benders decomposition. We evaluate the resulting investment plans through rolling-horizon Monte Carlo simulations with an out-of-sample analysis and demonstrate that the added robustness can significantly improve the operational reliability of the transition plans with a moderate increase in total cost.

Keywords: OR in energy, Clean energy transition, Multi-stage stochastic programming, Benders decomposition, Non-anticipativity relaxation

1. Introduction

Global sustainability goals and net-zero commitments are intensifying the need to decarbonize large-scale, energy-intensive facilities. Planning a clean energy transition for such facilities is difficult, as it requires transforming capital-intensive infrastructure while maintaining a reliable and continuous energy supply (Davis

*Corresponding author.

Email addresses: ahmet.sener@std.bogazici.edu.tr (Ahmet Emir Şener), burakkocuk@sabanciuniv.edu (Burak Kocuk), tugce.yuksel@sabanciuniv.edu (Tuğçe Yüksel)

et al., 2018). Long-term emission targets make the timing and mix of investments critical, yet these decisions must be made under uncertainty about how technology costs and efficiencies will evolve. These investments must also support energy management at a high temporal resolution, which is itself subject to operational uncertainties such as fluctuating demand and variable technology performance. In this paper, we study the clean energy transition planning problem of a campus-scale integrated electricity and heat system under uncertainty at both the strategic and operational levels.

Campus-scale clean energy transition planning is increasingly framed as an integrated infrastructure design problem that coordinates distributed generation, storage, and energy conversion technologies within interconnected multi-energy systems. Existing studies have examined carbon-neutral transitions, microgrid design, renewable energy and storage integration, and multi-energy system configurations (Mavromatidis & Petkov, 2021; Tian et al., 2022; Vergara-Zambrano et al., 2025), typically combining investment decisions with operational feasibility and infrastructure constraints. Many of these studies, however, rely on deterministic assumptions, which limits their ability to assess how uncertainty affects long-term system performance. Studies that do incorporate stochasticity tend to focus on energy prices, demand growth, renewable resource availability, and policy changes (Azimian et al., 2025; Fang et al., 2024; Guevara et al., 2024). Technology-driven cost reductions and efficiency improvements, by contrast, are treated only implicitly, despite their substantial influence on the timing and portfolio mix of renewable investment decisions. In addition, many of these studies adopt static planning formulations or rely heavily on representative days (Bacci et al., 2025; Huylo et al., 2025; Rathi & Zhang, 2022) and aggregated temporal resolutions (Guevara et al., 2024; Ioannou et al., 2019) to maintain computational tractability. Such simplifications can obscure the chronological dependencies needed to accurately capture storage dynamics, supply–demand balancing, and system feasibility, particularly under high renewable penetration (Chen et al., 2025; Gao et al., 2023; Marcy et al., 2022).

Solving models that jointly incorporate uncertainty and high-resolution operations is computationally challenging, because long-term investment decisions must remain feasible across many scenario paths. To address this difficulty, the literature offers a range of solution approaches for multi-stage stochastic mixed-integer programs. The extensive-form deterministic equivalent serves as a standard benchmark by reformulating the entire scenario tree into a single deterministic optimization problem (Birge & Louveaux, 1997). This approach, however, quickly becomes computationally prohibitive as the number of stages, scenarios, and operational periods grows. To mitigate this burden, sampling-based methods, such as scenario reduction (Dupačová et al., 2003), sample average approximation (Kleywegt et al., 2002), and stochastic decomposition (Higle & Sen, 1991), reduce or approximate the scenario representation. Nevertheless, they primarily address the size of the uncertainty model and leave unresolved the structural difficulties arising from integrality, interstage coupling, and non-anticipativity.

A complementary line of research instead exploits the problem structure through decomposition. Scenario decomposition (Mulvey & Ruszczyński, 1995) and progressive hedging (Rockafellar & Wets, 1991) decompose the problem along scenario paths and coordinate the resulting subproblems by penalizing violations of non-anticipativity, while nested Benders decomposition (Birge, 1985) recursively splits the scenario tree

into stagewise subproblems and approximates downstream recourse functions through Benders cuts. The application of these methods to problems with integer variables, however, often requires additional mechanisms such as outer branch-and-bound schemes or integer-valid cutting planes. Building on this line of work, stochastic dual dynamic programming (Pereira & Pinto, 1991) approximates future cost-to-go functions through Benders cuts generated iteratively over forward and backward passes of the scenario tree. Stochastic dual dynamic integer programming (Zou et al., 2019), together with its recent extensions (Füllner et al., 2024), adapts this framework to mixed-integer programs by employing integer-valid cuts such as Lagrangian cuts. These methods, however, typically rely on the stagewise independence of the underlying stochastic process, an assumption that does not hold in the setting considered in this paper, where technology costs and efficiencies evolve recursively across stages through advancement multipliers.

Motivated by the modeling gaps and methodological limitations identified above, we formulate a multi-stage stochastic programming (MSSP) model that jointly optimizes investment and operational decisions for renewable generation, storage, and heat-transfer technologies at a high temporal resolution. To capture strategic-level uncertainties in technology costs and efficiencies, we construct a scenario tree that represents the sequential revelation of technological progress over the planning horizon. We further introduce a robust reformulation based on box uncertainty sets to account for short-term operational uncertainties.

To solve the resulting large-scale model, we develop a Benders decomposition algorithm with partial non-anticipativity relaxation, grounded in scenario-based decomposition. The decision variables are split between the master problem and the subproblems according to the difficulty of restoring their associated non-anticipativity constraints. Specifically, investment variables and their non-anticipativity constraints are retained in the master problem, while operational variables are assigned to scenario-wise subproblems in which the corresponding non-anticipativity constraints are temporarily relaxed. Rather than progressively correcting non-anticipativity violations within each Benders iteration, we restore the non-anticipativity of operational decisions at the end of the algorithm by solving a smaller linear program (LP). Owing to the structural properties of the proposed model, this correction step is proven to recover a feasible solution with a bounded increase in the objective function value.

This paper makes three main contributions. (i) We formulate an MSSP model for the clean energy transition of a campus-scale integrated electricity–heat system under uncertain technological progress, and complement it with a robust reformulation that captures operational uncertainty. (ii) We develop a Benders decomposition algorithm with partial non-anticipativity relaxation with a correction step. We further strengthen the algorithm with valid inequalities and a two-phase cut-addition strategy. (iii) We evaluate the nominal and robust investment plans through rolling-horizon Monte Carlo simulations in an out-of-sample analysis over a broader range of operational realizations. Together, these developments enable us to solve a clean electricity–heat transition case study for the Middle East Technical University (METU) campus in Ankara, Turkey, comprising 9,435,001 constraints and 10,352,904 decision variables. In comparison, our earlier work (Şener et al., 2025), which did not employ the proposed method, could only address the clean electricity transition under uncertain technological advancement without operational uncertainty considerations.

The rest of the paper is organized as follows. Section 2 presents the MSSP model, its robust reformulation, and valid inequalities, while Section 3 characterizes the technological and operational uncertainties they address. Section 4 develops the proposed Benders decomposition algorithm and the out-of-sample evaluation procedure. Section 5 reports the computational experiments and the campus-scale case study. Finally, Section 6 concludes the paper.

2. Optimization Model

Table 1: Nomenclature

<i>Sets</i>		β_{real}	Real discount rate, calculated as $\beta_{\text{real}} = \frac{1+\beta_{\text{nominal}}}{1+\zeta} - 1$.
Ψ	Set of stages.	β	Discount factor, calculated as $\beta = \frac{1}{1+\beta_{\text{real}}}$.
$\mathcal{N}$	Set of nodes in the scenario tree, $\mathcal{N} = \{0, \dots, N\}$.	π_n	Probability of node $n \in \mathcal{N}$ in the scenario tree.
$\mathcal{N}_l$	Set of leaf nodes in the scenario tree.	ψ_n	Stage that node $n \in \mathcal{N}$ belongs to.
$\mathcal{T}$	Set of investment periods, $\mathcal{T} = \{0, \dots, T\}$.	$\mu_{n,t}$	Ancestor node of $n \in \mathcal{N}$ in period $t \in \mathcal{T}$.
$\mathcal{T}_n$	Subset of investment periods associated with node $n \in \mathcal{N}$.	$\kappa_{q,t,n}$	Sub-period, period, and node immediately preceding sub-period $q \in \mathcal{Q}$ in period $t \in \mathcal{T}_n$ at node $n \in \mathcal{N}$.
$\mathcal{Q}$	Set of operational sub-periods within each investment period, $\mathcal{Q} = \{1, \dots, Q\}$.	$\hat{\kappa}_{q,t,n}$	Period immediately preceding sub-period $q \in \mathcal{Q}$ in period $t \in \mathcal{T}_n$ at node $n \in \mathcal{N}$.
$\mathcal{R}$	Set of energy categories, $\mathcal{R} = \{\text{electricity, heat}\}$.	$\alpha_{u,n}$	Installation cost of technology $u \in \mathcal{U}$ at node $n \in \mathcal{N}$.
$\mathcal{J}$	Set of energy generation technologies.	$\hat{\alpha}_{u,t}$	Total O&M cost over the remaining planning horizon for technology $u \in \mathcal{U}$ installed in period $t \in \mathcal{T}$.
$\mathcal{S}$	Set of energy storage technologies.	$\lambda_{r,t}$	Cost of purchasing energy $r \in \mathcal{R}$ exogenously in period $t \in \mathcal{T}$.
$\mathcal{A}$	Set of heat transfer technologies.	$\delta_{r,q,t}$	Energy $r \in \mathcal{R}$ demand in sub-period $q \in \mathcal{Q}$ of period $t \in \mathcal{T}$.
$\mathcal{U}$	Set of all technologies, $\mathcal{U} = \mathcal{J} \cup \mathcal{S} \cup \mathcal{A}$.	$v_{s,r,t,t'}$	Energy $r \in \mathcal{R}$ storage capacity of technology $s \in \mathcal{S}$, during period $t' \in \mathcal{T}$ for installations made in period $t \in \mathcal{T}'_{s,[t']}$.
$\mathcal{T}_{u,[t]}$	Set of periods t' during which technology $u \in \mathcal{U}$, installed in period $t \in \mathcal{T}$, remains operational, i.e., $t \leq t' < t + \tau_{(u,t)}$.	φ_a	Heat transfer capacity of technology $a \in \mathcal{A}$.
$\mathcal{T}'_{u,[t']}$	Set of possible installation periods t of technology $u \in \mathcal{U}$ that is operational in period $t' \in \mathcal{T}$, i.e., $t \leq t' < t + \tau_{(u,t)}$.	$\sigma_{u,t,t'}$	Degradation rate of technology $u \in \mathcal{U}$, in period $t' \in \mathcal{T}$, for units installed in period $t \in \mathcal{T}'_{u,[t']}$.
<i>Decision variables</i>		η_r^+, η_r^-	Charging and discharging efficiencies of energy $r \in \mathcal{R}$ storage technologies, respectively.
$x_{u,t,n}$	Number of installations of technology $u \in \mathcal{U}$ in period $t \in \mathcal{T}_n$ at node $n \in \mathcal{N}$.	η_r	Self discharge rate of energy $r \in \mathcal{R}$ storage technologies.
$h_{r,q,t,n}$	Energy $r \in \mathcal{R}$ stored at the end of sub-period $q \in \mathcal{Q}$ in period $t \in \mathcal{T}_n$ at node $n \in \mathcal{N}$, carried forward to the next sub-period.	$\text{cop}_{a,q,t,n}$	Coefficient of performance in sub-period $q \in \mathcal{Q}$ for technology $a \in \mathcal{A}$ installed in period $t \in \mathcal{T}_n$ at node $n \in \mathcal{N}$.
$h_{r,q,t,n}^+$	Energy $r \in \mathcal{R}$ charged into storage during sub-period $q \in \mathcal{Q}$ in period $t \in \mathcal{T}_n$ at node $n \in \mathcal{N}$.	$\rho_{u,t}$	Spatial requirement of technology $u \in \mathcal{U}$, in period $t \in \mathcal{T}$.
$h_{r,q,t,n}^-$	Energy $r \in \mathcal{R}$ discharged from storage during sub-period $q \in \mathcal{Q}$ in period $t \in \mathcal{T}_n$ at node $n \in \mathcal{N}$.	$\tau_{u,t}$	Economic lifetime of technology $u \in \mathcal{U}$ installed in period $t \in \mathcal{T}$.
$p_{r,q,t,n}$	Energy $r \in \mathcal{R}$ purchased exogenously in sub-period $q \in \mathcal{Q}$ in period $t \in \mathcal{T}_n$ at node $n \in \mathcal{N}$.	u_u	Number of technology $u \in \mathcal{U}$ existing at the beginning of the planning horizon.
$o_{r,q,t,n}$	Energy $r \in \mathcal{R}$ used from inventory and generation in sub-period $q \in \mathcal{Q}$ of period $t \in \mathcal{T}_n$ at node $n \in \mathcal{N}$.	γ_r	Maximum fraction of annual demand for $r \in \mathcal{R}$ that can be met by emission-causing exogenous purchases in the final period.
$m_{a,q,t,t',n}$	Heat transferred in sub-period $q \in \mathcal{Q}$ in period $t \in \mathcal{T}_n$ at node $n \in \mathcal{N}$ by technology $a \in \mathcal{A}$ that was installed in period $t \in \mathcal{T}'_{a,[t']}$.	ϕ_t	Available budget for installations in period $t \in \mathcal{T}$.
<i>Parameters</i>		$\varkappa_t$	Cumulative maximum installation area available up to period $t \in \mathcal{T}$.
$\Gamma_{j,r,q,t,t',n}$	Energy $r \in \mathcal{R}$ generation of technology $j \in \mathcal{J}$, in sub-period $q \in \mathcal{Q}$ of period $t' \in \mathcal{T}_n$ at node $n \in \mathcal{N}$, for installations made in period $t \in \mathcal{T}'_{j,[t']}$.		
β_{nominal}	Nominal discount rate applied to the periods.		
ζ	Inflation rate applied to the periods.		

2.1. Multi-Stage Stochastic Programming Model

The deterministic equivalent of the proposed MSSP model is formulated using the notation defined in Table 1. Technology costs and efficiency-related performance profiles, including renewable energy generation and coefficients of performance, are indexed by the nodes in the scenario-tree $\mathcal{N}$ to capture their stage-wise realization under uncertainty. Each node $n \in \mathcal{N}$ comprises a set of periods $\mathcal{T}_n$, in which long-term technology investment decisions are made, and each period is further divided into sub-periods $\mathcal{Q}$, in which short-term operational decisions are made to meet energy demand. The resulting mixed-integer linear program given below jointly optimizes these investment and operational decisions for clean energy transition planning in an integrated electricity–heat system.

$$\text{minimize } \sum_{n \in \mathcal{N} \setminus \{0\}} \sum_{t \in \mathcal{T}_n} \pi_n \beta^{(t-1)} \left(\sum_{u \in \mathcal{U}} (\alpha_{u,n} + \hat{\alpha}_{u,t}) x_{u,t,n} + \sum_{r \in \mathcal{R}} \sum_{q \in \mathcal{Q}} \lambda_{r,t} p_{r,q,t,n} \right) \quad (1a)$$

subject to

$$o_{r,q,t',n} + \sum_{a \in \mathcal{A}} \sum_{t \in \mathcal{T}'_{a,[t']}} \mathbb{1}_{a,r,q,t,t',n} m_{a,q,t,t',n} \geq \delta_{r,q,t'} \quad n \in \mathcal{N} \setminus \{0\}, t' \in \mathcal{T}_n, q \in \mathcal{Q}, r \in \mathcal{R} \quad (1b)$$

$$p_{r,q,t',n} - h_{r,q,t',n}^+ + h_{r,q,t',n}^- + \sum_{j \in \mathcal{J}} \sum_{t \in \mathcal{T}'_{j,[t']}} \Gamma_{j,r,q,t,t',n} x_{j,t,\mu_{t,n}} \geq o_{r,q,t',n} \quad n \in \mathcal{N} \setminus \{0\}, t' \in \mathcal{T}_n, q \in \mathcal{Q}, r \in \mathcal{R} \quad (1c)$$

$$h_{r,q,t',n} = \eta_r h_{r,\kappa_{q,t',n}} + \eta_r^+ h_{r,q,t',n}^+ - \frac{h_{r,q,t',n}^-}{\eta_r} \quad n \in \mathcal{N} \setminus \{0\}, t' \in \mathcal{T}_n, q \in \mathcal{Q}, r \in \mathcal{R} \quad (1d)$$

$$h_{r,q,t',n} \leq \sum_{s \in \mathcal{S}} \sum_{t \in \mathcal{T}'_{s,[t']}} v_{s,r,t,t'} x_{s,t,\mu_{t,n}} \quad n \in \mathcal{N}, t' \in \mathcal{T}_n, q \in \mathcal{Q}, r \in \mathcal{R} \quad (1e)$$

$$m_{a,q,t,t',n} \leq \varphi_a x_{a,t,\mu_{t,n}} \quad n \in \mathcal{N}, t' \in \mathcal{T}_n, a \in \mathcal{A}, t \in \mathcal{T}'_{a,[t']}, q \in \mathcal{Q} \quad (1f)$$

$$\sum_{q \in \mathcal{Q}} p_{r,q,t,n} \leq \gamma_r \sum_{q \in \mathcal{Q}} \delta_{r,q,t} \quad n \in \mathcal{N}, r \in \mathcal{R} \quad (1g)$$

$$\sum_{u \in \mathcal{U}} \alpha_{u,n} x_{u,t,n} \leq \phi_t \quad n \in \mathcal{N} \setminus \{0\}, t \in \mathcal{T}_n \quad (1h)$$

$$\sum_{u \in \mathcal{U}} \sum_{t \in \mathcal{T}'_{u,[t']}} \rho_{u,t} x_{u,t,\mu_{t,n}} \leq \kappa_{t'} \quad n \in \mathcal{N} \setminus \{0\}, t' \in \mathcal{T}_n \quad (1i)$$

$$x_{u,0,0} = \iota_u \quad u \in \mathcal{U} \quad (1j)$$

$$\left\lfloor \frac{\phi_t}{\alpha_{u,n}} \right\rfloor \geq x_{u,t,n}, \quad x_{u,t,n} \in \mathbb{Z}_+ \quad n \in \mathcal{N}, t \in \mathcal{T}_n, u \in \mathcal{U} \quad (1k)$$

$$m_{a,q,t,t',n} \in \mathbb{R}_+ \quad n \in \mathcal{N}, t' \in \mathcal{T}_n, t \in \mathcal{T}'_{a,[t']}, q \in \mathcal{Q}, a \in \mathcal{A} \quad (1l)$$

$$p_{r,q,t,n}, h_{r,q,t,n}, h_{r,q,t,n}^+, h_{r,q,t,n}^-, o_{r,q,t,n} \in \mathbb{R}_+ \quad n \in \mathcal{N}, t \in \mathcal{T}_n, q \in \mathcal{Q}, r \in \mathcal{R}. \quad (1m)$$

The objective function (1a) minimizes the expected discounted total cost over the planning horizon. The cost components include technology installation costs, operation and maintenance (O&M) costs, and exogenous energy procurement costs. In each sub-period, energy demand must be satisfied through a combination of on-site generation, exogenous energy procurement, storage discharging, and heat transfer operations. We assume that energy supplied by heat transfer technologies cannot be charged into heat storage.

To impose this assumption, we introduce auxiliary decision variables $o_{r,q,t,n} \geq 0$ in the demand satisfaction constraint sets (1b) and (1c), whose non-negativity restricts storage charging to be supplied only by on-site generation and exogenous procurement. The indicator variables defined in (2) assign energy-category-specific coefficients to heat transfer operations, taking distinct values for heat and electricity, respectively.

$$\mathbb{I}_{a,r,q,t,t',n} = \begin{cases} 1 - \sigma_{a,t,t'}, & \text{if } r = \text{heat} \\ \frac{-1}{\text{cop}_{a,q,t,n}}, & \text{if } r = \text{electricity} \end{cases} \quad a \in \mathcal{A}, q \in \mathcal{Q}, n \in \mathcal{N}, t' \in \mathcal{T}_n, t \in \mathcal{T}'_{a,[t']} \quad (2)$$

Storage balance constraints (1d) govern storage dynamics by accounting for charging and discharging efficiencies as well as self-discharge losses. The stored energy levels and heat transfer capacities are bounded by the installed technology capacities through constraints (1e) and (1f), respectively. Notably, heat transfer decision variables and capacities are modeled in a disaggregated form to accurately capture technology degradation and the coefficient of performance.

To ensure environmental compliance, the percentage of annual demand that can be met by emission-causing exogenous energy purchases is constrained by (1g) for each energy category. Financial and spatial feasibility are enforced through budget constraints (1h), which cap annual installations based on available capital, and spatial constraints (1i), which limit the cumulative footprint of installed systems. The integration of pre-existing technologies at the start of the planning horizon is handled by initialization constraints (1j).

Installation decisions for energy technologies are represented by the non-negative integer variables defined in (1k). Their upper bounds, which follow from the available budget and the corresponding per-unit installation costs, are imposed explicitly as valid bounds on the integer variables. The operational decision variables are defined as non-negative continuous variables in (1l)–(1m).

2.2. Robust Reformulation

Alongside the investment plan obtained by solving model (1) with nominal operational profiles, we construct more conservative plans through a robust reformulation that protects the investment decisions against unfavorable realizations of the uncertain operational parameters: energy demand, renewable generation, and the coefficient of performance. We enclose these parameters in a box uncertainty set $[v_{d,i} - \epsilon \hat{v}_{d,i}, v_{d,i} + \epsilon \hat{v}_{d,i}]$, where the nominal value and its deviation are respectively denoted as $v_{d,i}$ and $\hat{v}_{d,i} \geq 0$, and the robustness parameter is $\epsilon \geq 0$. The uncertain parameters appear only in the demand satisfaction constraints (1b) and (1c), where each enters monotonically: feasibility is most tightly constrained when demand is high and when renewable generation and the coefficient of performance are low. The worst case over the uncertainty set is therefore attained at a vertex of the box, so it suffices to protect against this single worst-case realization rather than the entire set. Accordingly, we displace each parameter to the adverse end of its interval—energy demand to $v_{d,i} + \epsilon \hat{v}_{d,i}$, and renewable generation and the coefficient of performance to $v_{d,i} - \epsilon \hat{v}_{d,i}$ —and obtain the investment decisions under these displaced values. Setting $\epsilon = 0$ collapses the uncertainty set to the nominal point and recovers the original model, whereas larger values enlarge the set and yield increasingly conservative investment decisions. The robust model therefore retains the structure of the nominal one and can be solved in the same way.

2.3. Derivation of Valid Inequalities

In this subsection, we derive a family of valid inequalities that are expressed solely in terms of the investment decision variables.

Proposition 1. *For each $r \in \mathcal{R}$, the following inequalities are valid for the MSSP model (1) when $\gamma_r = 0$:*

$$\begin{aligned} & \eta_r^- \eta_r \sum_{s \in \mathcal{S}} \sum_{t \in \mathcal{T}_{s, [\hat{k}_{\underline{q}}, T, n]}'} v_{s, r, t, \hat{k}_{\underline{q}}, T, n} x_{s, t, \mu_{t, n}} + \sum_{q=\underline{q}}^{\bar{q}} \sum_{j \in \mathcal{J}} \sum_{t \in \mathcal{T}_{j, [T]}'} \Gamma_{j, r, q, t, T, n} x_{j, t, \mu_{t, n}} \\ & + \sum_{q=\underline{q}}^{\bar{q}} \sum_{a \in \mathcal{A}} \sum_{t \in \mathcal{T}_{a, [T]}'} s_{a, r, t, T} x_{a, t, \mu_{t, n}} \geq \sum_{q=\underline{q}}^{\bar{q}} \delta_{r, q, T} \end{aligned} \quad r \in \mathcal{R}, n \in \mathcal{N}_l, \underline{q} \in \mathcal{Q}, \bar{q} \in \{\underline{q}, \dots, Q\} \quad (3)$$

where the coefficients $s_{a, r, t, T}$ are defined as follows:

$$s_{a, r, t, T} = \begin{cases} \varphi_a (1 - \sigma_{a, t, T}), & \text{if } r = \text{heat} \\ 0, & \text{if } r = \text{electricity} \end{cases} \quad a \in \mathcal{A}, t \in \mathcal{T}_{a, [T]}' \quad (4)$$

Proof. Consider a leaf node $n \in \mathcal{N}_l$ and an energy category $r \in \mathcal{R}$. For $\gamma_r = 0$, the emission constraints (1g) prohibit exogenous energy purchases in period T . Consequently, the variables $p_{r, q, T, n}$ can be eliminated from (1c). Substituting the discharge variables $h_{r, q, T, n}^-$ using the storage balance constraints (1d) yields (5):

$$\eta_r^- \eta_r h_{r, \kappa_{\underline{q}}, T, n} - \eta_r^- h_{r, q, T, n} + (\eta_r^- \eta_r^+ - 1) h_{r, q, T, n}^+ + \sum_{j \in \mathcal{J}} \sum_{t \in \mathcal{T}_{j, [T]}'} \Gamma_{j, r, q, t, T, n} x_{j, t, \mu_{t, n}} \geq o_{r, q, T, n} \quad q \in \mathcal{Q} \quad (5)$$

Given that $\eta_r^+, \eta_r^- \in (0, 1)$, the coefficient $(\eta_r^- \eta_r^+ - 1)$ of $h_{r, q, T, n}^+$ in (5) is strictly negative. Since $h_{r, q, T, n}^+ \geq 0$ and the inequality is of the form “ $\geq$ ”, dropping this term preserves its validity, resulting in (6):

$$\eta_r^- (\eta_r h_{r, \kappa_{\underline{q}}, T, n} - h_{r, q, T, n}) + \sum_{j \in \mathcal{J}} \sum_{t \in \mathcal{T}_{j, [T]}'} \Gamma_{j, r, q, t, T, n} x_{j, t, \mu_{t, n}} \geq o_{r, q, T, n} \quad q \in \mathcal{Q} \quad (6)$$

Summing (6) over the interval $q \in \{\underline{q}, \dots, \bar{q}\}$ for any $\underline{q} \in \mathcal{Q}$ and $\bar{q} \in \{\underline{q}, \dots, Q\}$ yields (7):

$$\sum_{q=\underline{q}}^{\bar{q}} \left(\eta_r^- (\eta_r h_{r, \kappa_{\underline{q}}, T, n} - h_{r, q, T, n}) + \sum_{j \in \mathcal{J}} \sum_{t \in \mathcal{T}_{j, [T]}'} \Gamma_{j, r, q, t, T, n} x_{j, t, \mu_{t, n}} \right) \geq \sum_{q=\underline{q}}^{\bar{q}} o_{r, q, T, n} \quad \underline{q} \in \mathcal{Q}, \bar{q} \in \{\underline{q}, \dots, Q\} \quad (7)$$

Upon rearranging the terms, the interior storage variables $(h_{r, q, T, n})_{q=\underline{q}+1}^{\bar{q}-1}$ take the coefficient $\eta_r^- (\eta_r - 1)$, and $h_{r, \bar{q}, T, n}$ takes the coefficient $-\eta_r^-$. As these coefficients are negative and the storage variables are nonnegative, dropping these terms maintains the validity of the inequality, providing (8):

$$\eta_r^- \eta_r h_{r, \kappa_{\underline{q}}, T, n} + \sum_{q=\underline{q}}^{\bar{q}} \sum_{j \in \mathcal{J}} \sum_{t \in \mathcal{T}_{j, [T]}'} \Gamma_{j, r, q, t, T, n} x_{j, t, \mu_{t, n}} \geq \sum_{q=\underline{q}}^{\bar{q}} o_{r, q, T, n} \quad \underline{q} \in \mathcal{Q}, \bar{q} \in \{\underline{q}, \dots, Q\} \quad (8)$$

Next, multiplying the storage capacity constraints (1e) by $\eta_r^- \eta_r$ gives (9).

$$\eta_r^- \eta_r \sum_{s \in \mathcal{S}} \sum_{t \in \mathcal{T}_{s, [\hat{k}_{\underline{q}}, T, n]}'} v_{s, r, t, \hat{k}_{\underline{q}}, T, n} x_{s, t, \mu_{t, n}} \geq \eta_r^- \eta_r h_{r, \kappa_{\underline{q}}, T, n} \quad \underline{q} \in \mathcal{Q} \quad (9)$$

Combining (8) and (9) eliminates the remaining storage variables, yielding (10):

$$\eta_r \eta_r \sum_{s \in \mathcal{S}} \sum_{t \in \mathcal{T}'_{s, [\hat{k}_{\underline{q}, T, n}]}} v_{s, r, t, \hat{k}_{\underline{q}, T, n}} x_{s, t, \mu_{t, n}} + \sum_{q=\underline{q}}^{\bar{q}} \sum_{j \in \mathcal{J}} \sum_{t \in \mathcal{T}'_{j, [T]}} \Gamma_{j, r, q, t, T, n} x_{j, t, \mu_{t, n}} \geq \sum_{q=\underline{q}}^{\bar{q}} o_{r, q, T, n} \quad \underline{q} \in \mathcal{Q}, \bar{q} \in \{\underline{q}, \dots, Q\} \quad (10)$$

On the demand side, summing the demand satisfaction constraints (1b) over $q \in \{\underline{q}, \dots, \bar{q}\}$ provides (11):

$$\sum_{q=\underline{q}}^{\bar{q}} \left(o_{r, q, T, n} + \sum_{a \in \mathcal{A}} \sum_{t \in \mathcal{T}'_{a, [T]}} \mathbb{1}_{a, r, q, t, T, n} m_{a, q, t, T, n} \right) \geq \sum_{q=\underline{q}}^{\bar{q}} \delta_{r, q, T} \quad \underline{q} \in \mathcal{Q}, \bar{q} \in \{\underline{q}, \dots, Q\} \quad (11)$$

Since the coefficients $\mathbb{1}_{a, r, q, t, T, n}$ are negative when $r = \text{electricity}$, omitting these terms preserves the validity of the inequalities. To formalize this, we introduce new coefficients $\varsigma_{a, r, t, T}$ (defined in (4)) that retain only the nonnegative heat coefficients. Substituting these new coefficients and bounding the heat transfer variables by their capacity constraints (1f) yields the simplified inequalities (12):

$$\sum_{q=\underline{q}}^{\bar{q}} \left(o_{r, q, T, n} + \sum_{a \in \mathcal{A}} \sum_{t \in \mathcal{T}'_{a, [T]}} \varsigma_{a, r, t, T} x_{a, t, \mu_{t, n}} \right) \geq \sum_{q=\underline{q}}^{\bar{q}} \delta_{r, q, T} \quad \underline{q} \in \mathcal{Q}, \bar{q} \in \{\underline{q}, \dots, Q\} \quad (12)$$

Finally, combining (10) and (12) yields (3), showing that the family of inequalities given in (3) is valid. $\square$

Inequalities in (3) admit a direct physical interpretation. For any contiguous block of sub-periods in the final period, the aggregate demand within that block must be covered by the maximum energy that can be supplied by the installed resources. This amount consists of three components: (i) the usable energy recoverable from installed storage capacity, under the most favorable assumption that the storage system enters the block fully charged; (ii) the energy generated by installed technologies during the block; and, for heat demand only, (iii) the energy deliverable through installed heat-transfer capacity. Hence, these inequalities are necessary conditions for operational feasibility expressed solely in terms of the investment decision variables. Although inequalities in (3) are implied by (1) and therefore redundant for the original formulation, they strengthen the master problem in the Benders decomposition algorithm described in Section 4.1, which only contains investment decision variables. Section 4.1.1 describes how these inequalities are used as cutting planes within the proposed algorithm.

3. Uncertainty Modeling

This section describes how uncertainty is characterized and embedded in the proposed framework. We distinguish two categories: Technological uncertainty concerns the long-term evolution of technology installation costs and efficiencies over the planning horizon while operational uncertainty concerns the short-term variability of energy demand and technology performance within each operating period. Technological uncertainty determines the stochastic parameters in the MSSP model, while operational uncertainty is incorporated into the robust model extension and the out-of-sample Monte Carlo simulations.

3.1. Technological Uncertainty

We build the scenario tree in two steps, following our earlier work (Şener et al., 2025). First, we cluster historical cost-reduction and efficiency-improvement rates to estimate technological advancement multipliers

(Section 3.1.1). Second, we propagate these multipliers from an initial deterministic state to obtain the cost and performance parameters at every node of the scenario tree (Section 3.1.2).

3.1.1. Technological Advancement Multipliers

Motivated by Moore's Law, which characterizes technological performance improvement as exponential over time (Moore, 1998), we model technological progress through logarithmic improvement rates. Let $v_u^t \in \mathbb{R}_+^2$ denote the cost and efficiency of technology $u \in \mathcal{U}$ in year $t \in \{1, \dots, T\}$. For a stage length τ , the τ -year logarithmic improvement rates for cost ($\xi_{u,1}^t$) and efficiency ($\xi_{u,2}^t$) are defined as:

$$\xi_{u,1}^t := \ln \left(\frac{v_{u,1}^t}{v_{u,1}^{t+\tau}} \right), \quad \xi_{u,2}^t := \ln \left(\frac{v_{u,2}^{t+\tau}}{v_{u,2}^t} \right) \quad t = 1, \dots, T - \tau.$$

To capture the empirical distribution of these improvements, we partition the observed rate vectors $\xi_u^t = (\xi_{u,1}^t, \xi_{u,2}^t)^\top \in \mathbb{R}_+^2$ into $|K_u|$ clusters. We determine the optimal cluster representatives $w_u^k \in \mathbb{R}_+^2$ and the binary assignment variables $z_{u,t,k} \in \{0, 1\}$ by solving a mixed-integer nonlinear program (MINLP) that minimizes the within-cluster sum of squared ℓ_2 -distances:

$$\min_{w_u^k \in \mathbb{R}_+^2, z_{u,t,k} \in \{0,1\}} \left\{ \sum_{t=1}^{T-\tau} \sum_{k \in K_u} z_{u,t,k} \|w_u^k - \xi_u^t\|_2^2 : \sum_{k \in K_u} z_{u,t,k} = 1, t = 1, \dots, T - \tau \right\}.$$

Let $\tilde{w}_u^k$ and $\tilde{z}_{u,t,k}$ be the optimal solutions to this MINLP. We define the τ -year improvement multipliers for each cluster k as $\Xi_{u,1}^k := \exp(-\tilde{w}_{u,1}^k)$ and $\Xi_{u,2}^k := \exp(\tilde{w}_{u,2}^k)$. We then assume the technological parameters evolve according to the multiplier vector $\Xi_u^k = (\Xi_{u,1}^k, \Xi_{u,2}^k)^\top$ with probability $\chi_u^k := \frac{1}{T-\tau} \sum_{t=1}^{T-\tau} \tilde{z}_{u,t,k}$. This probability represents the proportion of observations assigned to cluster $k \in K_u$.

3.1.2. Scenario Tree Construction

Using the multipliers derived above, we construct the scenario tree following the standard finite, rooted-tree representation of the underlying stochastic process (Dupačová et al., 2000). An independent tree is generated for each technology $u \in \mathcal{U}$, where $T_u = (N(T_u), E(T_u))$ is a full, balanced tree with $\sum_{i=1}^{|\Psi|} |K_u|^{i-1}$ nodes. Each node $n \in N(T_u)$ carries three attributes: an installation cost, a node probability π_n , and a technology performance vector (e.g., energy generation output or coefficient of performance).

To illustrate the tree construction, consider an energy generation technology $j \in \mathcal{J}$. The root node $n_{\text{root}} \in N(T_j)$ represents the deterministic initial state at stage 1 (i.e., $\psi_{n_{\text{root}}} = 1$), where the initial installation cost $\alpha_{j,1}$ and the initial energy generation profile $(\Gamma_{j,r,q,t,1})_{t \in \mathcal{T}_1}$ are known, deterministic parameters. The root node probability is initialized as $\pi_{n_{\text{root}}} = 1$.

Each non-leaf node branches into $|K_j|$ child nodes, each corresponding to the realization of a distinct cluster $k \in K_j$. For any non-root node $n \in N(T_j) \setminus \{n_{\text{root}}\}$, let $p(n)$ denote its immediate parent and k_n the cluster realization governing the transition from $p(n)$ to n . The edge set is then defined as $E(T_j) = \{(p(n), n) : n \in N(T_j) \setminus \{n_{\text{root}}\}\}$. The installation cost, node probability, and generation vector are then computed recursively as:

$$\alpha_{j,n} = \alpha_{j,p(n)} \Xi_{j,1}^{k_n}, \quad \pi_n = \pi_{p(n)} \chi_j^{k_n} \quad j \in \mathcal{J}, n \in N(T_j) \setminus \{n_{\text{root}}\}$$

$$\Gamma_{j,r,q,t,t,n} = \Gamma_{j,r,q,\bar{t},\bar{p}(n)} \Xi_{j,2}^{k_n} \quad j \in \mathcal{J}, r \in \mathcal{R}, q \in \mathcal{Q}, n \in N(T_j) \setminus \{n_{\text{root}}\}, t \in \mathcal{T}_n, \bar{t} \in \mathcal{T}_{p(n)}.$$

For any period $t' > t$, the generation profile is further adjusted by a degradation factor $\sigma_{j,t,t'}$:

$$\Gamma_{j,r,q,t,t',n} = \Gamma_{j,r,q,t,t,n} \cdot (1 - \sigma_{j,t,t'}) \quad j \in \mathcal{J}, r \in \mathcal{R}, q \in \mathcal{Q}, n \in N(T_j), t \in \mathcal{T}_n, t' \in \mathcal{T}_{j,[t]}.$$

Once the individual trees T_u have been built for all $u \in \mathcal{U}$, they are assembled into a unified scenario tree $T = (N(T), E(T))$. Let $N_s(T_u) := \{n_u \in N(T_u) : \psi_{n_u} = s\}$ denote the set of nodes in tree T_u belonging to stage $s \in \Psi$. Since the technologies are assumed to evolve independently, the node set at any stage s is the Cartesian product of the per-technology node sets $N_s(T_u)$ across all $u \in \mathcal{U}$. We further introduce a deterministic initialization node (node 0), the parent of n_{root} , which encodes the technologies already in place prior to the first decision stage and has probability $\pi_0 = 1$. Consequently, each node $n \in N(T)$ takes the form of a tuple $n = (n_u)_{u \in \mathcal{U}}$, capturing the joint realization of all technologies at that stage. The full node and edge sets of T are thus:

$$N(T) = \{0\} \cup \bigcup_{s \in \Psi} \prod_{u \in \mathcal{U}} N_s(T_u), \quad E(T) = \{(0, n_{\text{root}})\} \cup \{(m, n) \in N(T) \times N(T) : (m_u, n_u) \in E(T_u), u \in \mathcal{U}\}.$$

The unconditional probability of each node $n \in N(T)$ is obtained as the product of its component node probabilities, $\pi_n = \prod_{u \in \mathcal{U}} \pi_{n_u}$. The attributes encoded in the resulting scenario tree constitute the stochastic parameters of the MSSP model given in (1).

3.2. Operational Uncertainty

Energy demand, renewable generation, and the coefficient of performance of heat-transfer technologies are inherently uncertain, driven by exogenous factors such as weather conditions and human activity (Zakaria et al., 2020). We characterize these uncertainties through probability distributions derived from historical operational data via time-series clustering—a method commonly applied in clean energy transition planning to identify representative days (Li et al., 2022; Teichgraeber & Brandt, 2019). To this end, we divide a given historical time series into consecutive daily profiles. Let $\mathcal{D}$ denote the set of days and $v_{d,i}$ the observation for day $d \in \mathcal{D}$ at hour $i \in \{1, \dots, 24\}$. To cluster daily profiles based on their temporal shape rather than absolute magnitude, we standardize each profile using its mean $\bar{v}_d = \frac{1}{24} \sum_{i=1}^{24} v_{d,i}$ and standard deviation s_d , where $s_d^2 = \frac{1}{24} \sum_{i=1}^{24} (v_{d,i} - \bar{v}_d)^2$. The standardized profile is then $g_{d,i} = (v_{d,i} - \bar{v}_d)/s_d$.

Given these standardized daily profiles, we partition them into clusters. Since no clustering method is universally superior (Li et al., 2022; Teichgraeber & Brandt, 2019), we evaluate several approaches—including k -means and hierarchical clustering under various parameter settings—and select the configuration with the lowest intra-cluster variance. This configuration is agglomerative hierarchical clustering with Euclidean distance and Ward's minimum-variance linkage criterion, restricted to at most five clusters.

Let $\mathcal{C}$ denote the set of clusters, and for each $c \in \mathcal{C}$, let $\mathcal{D}_c \subseteq \mathcal{D}$ denote the subset of days assigned to it. The hourly cluster mean and variance are then computed as $\bar{g}_{c,i} = \frac{1}{|\mathcal{D}_c|} \sum_{d \in \mathcal{D}_c} g_{d,i}$ and $s_{c,i}^2 = \frac{1}{|\mathcal{D}_c|} \sum_{d \in \mathcal{D}_c} (g_{d,i} - \bar{g}_{c,i})^2$. Figure 1 illustrates the resulting clusters for the 2023 hourly electricity demand series of the case study introduced in Section 5.1, together with the representative profile of each cluster

(dashed lines). The representative profile of a cluster is the member profile with the minimum total Euclidean distance to all others in the cluster. These cluster statistics define the probability distributions used to represent operational uncertainty. For each cluster $c \in \mathcal{C}$, day $d \in \mathcal{D}_c$, and hour i , we sample the standardized value as $\tilde{g}_{d,i} \sim \mathcal{N}(g_{d,i}, s_{c,i}^2)$, introducing cluster variability at each hour around each observed value. Transforming back to the original scale gives $\tilde{v}_{d,i} = s_d \tilde{g}_{d,i} + \bar{v}_d \sim \mathcal{N}(v_{d,i}, (s_d s_{c,i})^2)$. Each resulting distribution is therefore centered on the observed value, with variance jointly governed by the within-day standard deviation of the original profile and the hour-specific variability of the corresponding cluster. Applying this procedure independently to each series, we obtain the probability distributions for the demand parameters $\delta_{r,q,t}$, the generation parameters $\Gamma_{j,r,q,t,t',n}$, and the coefficients of performance $\text{cop}_{a,q,t,n}$.




Figure 1: Clusters of standardized daily profiles of the 2023 hourly electricity demand series for the case study introduced in Section 5.1. Dashed lines show the representative profile of each cluster.

4. Methodology

In this section, we present our solution approach for the clean energy transition planning problem introduced in Section 2. We begin by recasting the model in a compact form and introducing a scenario-wise reformulation, which motivates the proposed decomposition scheme. Section 4.1 then develops the Benders decomposition algorithm with partial non-anticipativity relaxation, and Section 4.2 describes the rolling-horizon simulation used to evaluate the resulting investment plans in an out-of-sample analysis.

Let x and y denote the vectors of investment and operational decision variables, respectively. Let c and d represent their corresponding objective coefficient vectors. We define the constraint sets that do not couple investment and operational decision variables respectively as $\mathcal{X} := \{x : (1h), (1i), (1j), (1k)\}$ and $\mathcal{Y} := \{y : (1b), (1d), (1g), (1l), (1m)\}$. The constraints coupling these decisions—namely (1c), (1e), and (1f)—are expressed generically as $Ax + By \geq b$. Then, the model (1) can be provided in compact form as

$$\min_{x \in \mathcal{X}, y \in \mathcal{Y}} \{c^\top x + d^\top y : Ax + By \geq b\}. \quad (13)$$

We now provide a scenario-wise reformulation of the model with partial non-anticipativity relaxation. Let Ω denote the set of scenario paths in the scenario tree, and let Π_ω denote the probability of path $\omega \in \Omega$. For each path, let x_ω and y_ω denote the associated investment and operational decision vectors, with c_ω and d_ω the corresponding cost coefficient vectors (excluding node probabilities), and let A_ω , B_ω , and b_ω denote the path-specific constraint parameters. In a standard MSSP formulation, non-anticipativity constraints require that any two scenario paths passing through the same node take identical decisions at that node. We

relax these constraints for the operational variables while preserving them for the investment variables. The relaxation renders the operational feasible region separable across paths, yielding path-specific sets $\mathcal{Y}_\omega$. Then the relaxed problem becomes:

$$\min_{(x_\omega, y_\omega)_{\omega \in \Omega}} \left\{ \sum_{\omega \in \Omega} \Pi_\omega(c_\omega^\top x_\omega + d_\omega^\top y_\omega) : \begin{array}{ll} A_\omega x_\omega + B_\omega y_\omega \geq b_\omega, & \omega \in \Omega \\ y_\omega \in \mathcal{Y}_\omega, & \omega \in \Omega \\ (x_\omega)_{\omega \in \Omega} \in \mathcal{X} \end{array} \right\}. \quad (14)$$

Because non-anticipativity is no longer enforced on the operational variables, an optimal solution to (14) may assign different operational decisions to scenario paths that share the same node. Once an optimal solution to (14) is obtained, we restore non-anticipativity through the procedure described in Section 4.1.2.

4.1. Benders Decomposition with Partial Non-anticipativity Relaxation

For a fixed investment decision vector $(x_\omega^*)_{\omega \in \Omega}$, problem (14) decomposes by scenario path into $|\Omega|$ independent subproblems. For each scenario path $\omega \in \Omega$, the corresponding subproblem is given by

$$\mathbf{SP}_\omega(x_\omega^*) : \min_{y_\omega \in \mathcal{Y}_\omega} \{d_\omega^\top y_\omega : B_\omega y_\omega \geq b_\omega - A_\omega x_\omega^*\}, \quad \omega \in \Omega. \quad (15)$$

To approximate the objective value of each scenario-path subproblem within the master problem, we introduce auxiliary variables $(\theta_\omega)_{\omega \in \Omega}$ and progressively refine this approximation with Benders cuts generated over a sequence of iterations indexed by k . At iteration k , the scenario-path subproblems $\mathbf{SP}_\omega(x_\omega^{(k)})$ are solved at the current candidate investment vector $(x_\omega^{(k)})_{\omega \in \Omega}$. A feasible subproblem yields an optimal dual extreme point $\lambda_\omega^{(k)}$ that generates an optimality cut, whereas an infeasible subproblem yields a dual extreme ray $\mu_\omega^{(k)}$ that generates a feasibility cut. A newly generated point or ray is retained only if its associated cut is violated by the current candidate solution. Let $\Lambda_\omega^{\text{opt}}$ and $\Lambda_\omega^{\text{feas}}$ denote the sets of dual extreme points and extreme rays retained up to iteration k , respectively. The associated set of Benders cuts is then given by

$$C := \left\{ \begin{array}{ll} \theta_\omega \geq (\lambda_\omega)^\top (b_\omega - A_\omega x_\omega), & \omega \in \Omega, \lambda_\omega \in \Lambda_\omega^{\text{opt}}, \\ 0 \geq (\mu_\omega)^\top (b_\omega - A_\omega x_\omega), & \omega \in \Omega, \mu_\omega \in \Lambda_\omega^{\text{feas}} \end{array} \right\}. \quad (16)$$

Each element of C is a linear inequality in the master variables $(x_\omega, \theta_\omega)_{\omega \in \Omega}$. Using this cut set, the master problem is defined as

$$\mathbf{MP}(C) : \min_{(x_\omega, \theta_\omega)_{\omega \in \Omega}} \left\{ \sum_{\omega \in \Omega} \Pi_\omega(c_\omega^\top x_\omega + \theta_\omega) : (x_\omega)_{\omega \in \Omega} \in \mathcal{X}, (x_\omega, \theta_\omega)_{\omega \in \Omega} \text{ satisfies } C \right\}. \quad (17)$$

Algorithm 1 summarizes the proposed solution procedure, in which the scenario-path subproblems are solved in parallel at each iteration. In addition to the Benders cuts generated from these subproblems, the master problem is strengthened by valid inequalities: the `SEPARATIONKADANE` procedure described in Section 4.1.1 identifies violated valid inequalities (line 12), which are then added to the master problem together with the Benders cuts.

The procedure is organized into two phases that differ in how the integrality constraints and the cut generation are handled. In Phase 1, the integrality constraints of the master problem are relaxed, so that only

its linear relaxation is solved at each Benders iteration, and both feasibility and optimality cuts are added at every iteration to accelerate convergence. The algorithm switches to Phase 2 once the convergence of the relaxed master problem slows down, that is, when the Benders optimality tolerance is attained by the linear relaxation or the lower bound stagnates over a window of consecutive iterations. Phase 2 reinstates the integrality constraints; both cut types are added in its early iterations, whereas in the later iterations feasibility cuts are prioritized and optimality cuts are added only when all subproblems are feasible, so that the search concentrates on restoring subproblem feasibility. To limit the growth of the master problem, we remove cuts that are no longer binding: a cut whose slack ranks in the highest 50% of all cut slacks for 100 consecutive iterations is moved from C to a pool (line 24), and any cut in the pool that is violated by the current master solution is reintroduced to C . The algorithm terminates when the gap between the best lower bound from the master problem and the incumbent objective value falls within the prescribed Benders optimality tolerance (ε_B), after which a final step restores non-anticipativity of the operational decisions (line 29), as described in Section 4.1.2.

Algorithm 1 Benders decomposition framework for the MSSP problem

Input: Set of scenario paths Ω with their probabilities $(\Pi_\omega)_{\omega \in \Omega}$; model inputs $\mathcal{X}$ and $(c_\omega, d_\omega, b_\omega, A_\omega, B_\omega, \mathcal{Y}_\omega)_{\omega \in \Omega}$; functions $V_{\omega,r} : \{(x_\omega, q, \bar{q}) \in \mathcal{X}_\omega \times \mathcal{Q}^2 \mid q \leq \bar{q}\} \rightarrow \mathbb{R}$ for all $\omega \in \Omega$ and all $r \in \mathcal{R}$; Benders tolerance $\varepsilon_B \geq 0$, violation tolerance $\varepsilon_V \geq 0$.

Output: Solution x^*, y^* and objective value UB.

- 1: $LB^{(0)} \leftarrow -\infty$, $UB \leftarrow +\infty$, $k \leftarrow 1$, $C \leftarrow \emptyset$, Phase $\leftarrow 1$
- 2: Relax the integrality constraints of $MP(C)$.
- 3: **while** Phase = 1 **or** $|UB - LB^{(k-1)}|/|UB| > \varepsilon_B$ **do**
- 4: Solve problem $MP(C)$ and let $(x^{(k)}, \theta^{(k)})$ and $z_{lb}^{(k)}$ denote the optimal solution and lower bound, respectively.
- 5: $C_{\text{feas}} \leftarrow \emptyset$, $C_{\text{opt}} \leftarrow \emptyset$
- 6: **for all** $\omega \in \Omega$ **in parallel do**
- 7: Solve subproblem $SP_\omega(x_\omega^{(k)})$.
- 8: **if** $SP_\omega(x_\omega^{(k)})$ is infeasible **then**
- 9: Let $\mu_\omega^{(k)}$ denote a dual extreme ray associated with $SP_\omega(x_\omega^{(k)})$, $z_\omega \leftarrow +\infty$.
- 10: **if** $(\mu_\omega^{(k)})^\top (b_\omega - A_\omega x_\omega^{(k)}) \geq \varepsilon_V$ **then**
- 11: $C_{\text{feas}} \leftarrow C_{\text{feas}} \cup \{0 \geq (\mu_\omega^{(k)})^\top (b_\omega - A_\omega x_\omega)\}$
- 12: $C_{\text{feas}} \leftarrow C_{\text{feas}} \cup \text{SEPARATIONKADANE}(V_{\omega,r}, x_\omega^{(k)}, \mathcal{R}, \mathcal{Q})$
- 13: **else**
- 14: Let $\lambda_\omega^{(k)}$ and z_ω denote an optimal dual extreme point and the optimal objective value of $SP_\omega(x_\omega^{(k)})$, respectively.
- 15: **if** $(\lambda_\omega^{(k)})^\top (b_\omega - A_\omega x_\omega^{(k)}) - \theta_\omega^{(k)} \geq \varepsilon_V$ **then**
- 16: $C_{\text{opt}} \leftarrow C_{\text{opt}} \cup \{\theta_\omega \geq (\lambda_\omega^{(k)})^\top (b_\omega - A_\omega x_\omega)\}$
- 17: **if** Phase = 1 **or** in the first 200 Phase 2 iterations **or** an incumbent was found within the last 25 iterations **or** $C_{\text{feas}} = \emptyset$ **then**
- 18: $C \leftarrow C \cup C_{\text{opt}} \cup C_{\text{feas}}$
- 19: **else**
- 20: $C \leftarrow C \cup C_{\text{feas}}$
- 21: $z^{(k)} \leftarrow \sum_{\omega \in \Omega} \Pi_\omega (c_\omega^\top x_\omega^{(k)} + z_\omega)$
- 22: **if** $z^{(k)} < UB$ **then**
- 23: $UB \leftarrow z^{(k)}$, $x^* \leftarrow x^{(k)}$
- 24: Update C by pruning selected cuts and reinstating violated cuts from the pruned cuts pool.
- 25: $LB^{(k)} \leftarrow \max\{z_{lb}^{(k)}, LB^{(k-1)}\}$
- 26: **if** Phase = 1 **and** $(|UB - LB^{(k)}|/|UB| \leq \varepsilon_B$ **or** $|LB^{(k)} - LB^{(\max\{k-50, 0\})}|/|LB^{(k)}| \leq 10^{-3})$ **then**
- 27: Reintroduce the integrality constraints of $MP(C)$, Phase $\leftarrow 2$, $UB \leftarrow +\infty$.
- 28: $k \leftarrow k + 1$
- 29: Let y^* and z_{na} be the optimal solution and objective value of the correction LP described in Section 4.1.2.
- 30: **return** $x^*, y^*, z_{\text{na}} + \sum_{\omega \in \Omega} \Pi_\omega c_\omega^\top x_\omega^*$

4.1.1. Separation of Valid Inequalities

To accelerate the convergence of the Benders decomposition algorithm, we incorporate the valid inequalities in (3) in the master problem. By projecting necessary operational-feasibility conditions onto the investment space, these inequalities enable the master problem to anticipate operational constraints. These cuts thereby tighten the feasible region of the master problem by eliminating operationally infeasible investment decisions, reducing the number of Benders iterations required for convergence.

Adding all valid inequalities to the master problem is computationally impractical, as one such inequality is defined for each combination of a scenario path, an energy category, and a sub-period pair. To avoid incorporating this large set of inequalities upfront, we employ a separation procedure within the Benders algorithm. At each iteration, given the current master problem solution $x^{(k)}$, we solve a separation problem for each energy category and infeasible subproblem to identify the most violated inequalities, which are then added to the master problem alongside the Benders cuts.

Let $\mathcal{X}_\omega$ denote the projection of $\mathcal{X}$ onto the investment variables on the nodes of path ω . For a given scenario path $\omega \in \Omega$ and energy category $r \in \mathcal{R}$, we quantify the violation of inequality (3) at the current master problem solution $x_\omega^{(k)} \in \mathcal{X}_\omega$ using the function:

$$\begin{aligned} V_{\omega,r}(x_\omega^{(k)}, \underline{q}, \bar{q}) := & \eta_r \eta_r \sum_{s \in \mathcal{S}} \sum_{t \in \mathcal{T}'_{s, [\hat{k}_{\underline{q}, T, n_\omega}]}} v_{s,r,t, \hat{k}_{\underline{q}, T, n_\omega}} x_{s,t, \mu_t, n_\omega}^{(k)} + \sum_{q=\underline{q}}^{\bar{q}} \sum_{j \in \mathcal{J}} \sum_{t \in \mathcal{T}'_{j, [T]}} \Gamma_{j,r,q,t,T, n_\omega} x_{j,t, \mu_t, n_\omega}^{(k)} \\ & + \sum_{q=\underline{q}}^{\bar{q}} \sum_{a \in \mathcal{A}} \sum_{t \in \mathcal{T}'_{a, [T]}} \varsigma_{a,r,t,T} x_{a,t, \mu_t, n_\omega}^{(k)} - \sum_{q=\underline{q}}^{\bar{q}} \delta_{r,q,T} \end{aligned} \quad (18)$$

where n_ω denotes the leaf node associated with scenario path ω . The separation problem then seeks the sub-period pair that maximizes this violation (i.e., minimizes the function value):

$$\min_{1 \leq \underline{q} \leq \bar{q} \leq Q} V_{\omega,r}(x_\omega^{(k)}, \underline{q}, \bar{q}) \quad (19)$$

If the optimal value of (19) is strictly negative, the corresponding inequality (3) is violated at the current solution and therefore added to the master problem as a cutting plane. Otherwise, no violated inequality of this form exists for the given scenario path and energy category, and no cut is generated.

To solve (19), we construct an array of length $|Q|$ whose q -th entry stores the value $V_{\omega,r}(x_\omega^{(k)}, q, q)$, the violation associated with sub-period $q \in Q$. Identifying the most violated sub-period pair $(\underline{q}, \bar{q})$ then reduces to finding a contiguous subarray with the minimum sum. As established in Proposition 2, Kadane's algorithm (Bentley, 1984) solves this problem in linear time.

Proposition 2. *For a sequence $A = (a_1, \dots, a_n) \in \mathbb{R}^n$, Kadane's algorithm identifies a contiguous subarray with the minimum sum in $\mathcal{O}(n)$ time.*

Proof. Correctness and linear time complexity follow directly from the standard recurrence relation underlying Kadane's algorithm; for a detailed discussion, see Gries (1982). $\square$

For a given scenario path, solving (19) across all energy categories and separating the most violated

inequality requires $\mathcal{O}(|\mathcal{R}| \cdot |\mathcal{Q}|)$ time. Algorithm 2 details the complete procedure and generates the corresponding valid cuts.

Algorithm 2 Separation of valid inequalities using Kadane’s algorithm

Input: Functions $V_{\omega,r} : \{(x_\omega, \underline{q}, \bar{q}) \in \mathcal{X}_\omega \times \mathcal{Q}^2 \mid \underline{q} \leq \bar{q}\} \rightarrow \mathbb{R}$ for $\forall r \in \mathcal{R}$ with fixed $\omega \in \Omega, x_\omega^{(k)}, \mathcal{R}, \mathcal{Q}$, and tolerance $\varepsilon_V \geq 0$.
Output: Set of violated valid inequalities C_V .

```

1: function SEPARATIONKADANE( $V_{\omega,r}, x_\omega^{(k)}, \mathcal{R}, \mathcal{Q}$ )
2:    $C_V \leftarrow \emptyset$ 
3:   for all  $r \in \mathcal{R}$  do
4:      $\text{currentSum} \leftarrow V_{\omega,r}(x_\omega^{(k)}, 1, 1), \underline{q}_{\text{curr}} \leftarrow 1$  ▷ Local minimum sum ending at  $q$  and its start index
5:      $\text{bestSum} \leftarrow \text{currentSum}, \underline{q}^* \leftarrow 1, \bar{q}^* \leftarrow 1$  ▷ Global minimum sum over  $[1, q]$  and its boundaries
6:     for all  $q \in \{2, \dots, Q\}$  do
7:       if  $\text{currentSum} > 0$  then
8:          $\text{currentSum} \leftarrow V_{\omega,r}(x_\omega^{(k)}, q, q), \underline{q}_{\text{curr}} \leftarrow q$  ▷ Initiate a new locally optimal subarray starting at  $q$ 
9:       else
10:         $\text{currentSum} \leftarrow \text{currentSum} + V_{\omega,r}(x_\omega^{(k)}, q, q)$  ▷ Extend the previous subarray
11:       if  $\text{currentSum} < \text{bestSum}$  then
12:         $\text{bestSum} \leftarrow \text{currentSum}, \underline{q}^* \leftarrow \underline{q}_{\text{curr}}, \bar{q}^* \leftarrow q$  ▷ Update global minimum sum
13:     if  $\text{bestSum} < -\varepsilon_V$  then
14:        $C_V \leftarrow C_V \cup \{V_{\omega,r}(x_\omega, \underline{q}^*, \bar{q}^*) \geq 0\}$ 
15:   return  $C_V$ 

```

4.1.2. Restoring Non-anticipativity

Upon termination, Algorithm 1 returns an incumbent (x^*, y^*) of the relaxed problem (14), where x^* is non-anticipative by construction but y^* need not be. We restore non-anticipativity in a correction step that fixes x at x^* and re-optimizes y by solving an LP subject to the constraints of the original model (1) and the non-anticipativity constraints. We first show that, under a mild assumption on the stage structure, such an incumbent always admits a non-anticipative solution at x^* that is feasible for the original model and whose objective function value exceeds that of (x^*, y^*) by a bounded amount. To this end, let $\underline{t}_n$ and $\bar{t}_n$ denote the first and last periods of node $n \in \mathcal{N}$, and let $\bar{\mu}_n$ denote the set of immediate children of n in the scenario tree.

Proposition 3. *Let $(x_\omega^*, y_\omega^*)_{\omega \in \Omega}$ be an incumbent of the non-anticipativity-relaxed problem (14), and assume that each node in the final stage consists of more than one period. Then, there exists a non-anticipative solution $(x_\omega^*, \hat{y}_\omega)_{\omega \in \Omega}$ of (1) whose objective function value exceeds that of $(x_\omega^*, y_\omega^*)_{\omega \in \Omega}$ by at most*

$$\sum_{r \in \mathcal{R}} \sum_{n \in \mathcal{N} \setminus \mathcal{N}_I} \sum_{n' \in \bar{\mu}_n} \frac{\pi_{n'} \beta^{(\underline{t}_{n'}-1)} \lambda_{r, \underline{t}_{n'}} (\Delta_{r, n, n'})^+}{\eta_r^+}, \quad (20)$$

where $(\cdot)^+ := \max\{0, \cdot\}$ and $\Delta_{r, n, n'}$ is the residual of the inter-node storage-balance constraints (1d), evaluated at the node-wise solution $\tilde{y}$ obtained when these constraints are relaxed:

$$\Delta_{r, n, n'} := \tilde{h}_{r, 1, \underline{t}_{n'}, n'} - \eta_r \tilde{h}_{r, Q, \bar{t}_n, n} - \eta_r^+ \tilde{h}_{r, 1, \underline{t}_{n'}, n'}^+ + \frac{\tilde{h}_{r, 1, \underline{t}_{n'}, n'}^-}{\eta_r^-}. \quad (21)$$

Proof. Within $\mathbf{SP}_\omega(x_\omega^*)$, consecutive nodes along path ω are coupled through the storage-balance constraints (1d)—one per energy category—which relate the terminal sub-period of each node to the first

sub-period of its successor. Relaxing these inter-node constraints decomposes $\mathbf{SP}_\omega(x_\omega^*)$ into node-wise subproblems $\widetilde{\mathbf{SP}}_{\omega,i}(x_\omega^*)$, one per stage $i \in \Psi$, while retaining the within-node storage balances.

Consider two scenario paths $\omega_1, \omega_2 \in \Omega$ that pass through a common node $n \in \mathcal{N}$. The node-wise subproblems $\widetilde{\mathbf{SP}}_{\omega_1, \psi_n}(x_{\omega_1}^*)$ and $\widetilde{\mathbf{SP}}_{\omega_2, \psi_n}(x_{\omega_2}^*)$ have identical feasible regions and objective coefficients: their investment decisions agree because x^* is non-anticipative on all shared nodes, and their stochastic parameters agree because both are evaluated at n . The node-wise subproblem at n is therefore independent of the path containing n , so we may discard the duplicates and associate a single subproblem $\widetilde{\mathbf{SP}}_n(x_n)$ with each node $n \in \mathcal{N}$, where x_n is the common investment decision at n and its ancestors. Any non-anticipativity violation in y^* at a shared node is thus attributable solely to the relaxed inter-node storage-balance constraints.

Each $\widetilde{\mathbf{SP}}_n(x_n)$ is feasible: it is obtained by relaxing the inter-node constraints of $\mathbf{SP}_\omega(x_\omega^*)$, which is itself feasible because $(x_\omega^*, y_\omega^*)_{\omega \in \Omega}$ is feasible for (14). Let $\tilde{y}$ denote the non-anticipative solution obtained by solving these node-wise subproblems; it satisfies every constraint of (1) except, possibly, the relaxed inter-node storage-balance constraints. Because each $\widetilde{\mathbf{SP}}_n(x_n)$ relaxes the inter-node coupling of the corresponding per-path subproblem, the restriction of y^* to node n is feasible for $\widetilde{\mathbf{SP}}_n(x_n)$; the node-wise optimum at n therefore costs no more than y^* does at n , and summing over all nodes shows that the objective of $\tilde{y}$ is no greater than that of y^* .

It remains to reinstate the inter-node constraints while preserving non-anticipativity and feasibility. Consider a non-leaf node $n \in \mathcal{N} \setminus \mathcal{N}_l$ and a child $n' \in \bar{\mu}_n$. The reinstated constraint couples the terminal storage $h_{r,Q,\bar{t}_n,n}$ of n to the first sub-period of n' through

$$h_{r,1,\underline{t}_{n'},n'} = \eta_r h_{r,Q,\bar{t}_n,n} + \eta_r^+ h_{r,1,\underline{t}_{n'},n'}^+ - \frac{h_{r,1,\underline{t}_{n'},n'}^-}{\eta_r^-}. \quad (22)$$

Evaluated at the node-wise solution $\tilde{y}$, this equation need not hold; its residual $\Delta_{r,n,n'}$ is defined as in (21). A positive residual indicates that the first sub-period of n' draws more stored energy than its inputs supply. We restore the balance at n' itself, leaving the storage carried over from n untouched: we raise the charging from $\tilde{h}_{r,1,\underline{t}_{n'},n'}^+$ to $\hat{h}_{r,1,\underline{t}_{n'},n'}^+ := \tilde{h}_{r,1,\underline{t}_{n'},n'}^+ + \Delta_{r,n,n'}/\eta_r^+$ and meet it with an equal increase in the exogenous purchase, $\hat{p}_{r,1,\underline{t}_{n'},n'} := \tilde{p}_{r,1,\underline{t}_{n'},n'} + \Delta_{r,n,n'}/\eta_r^+$, in the first sub-period of n' . Because the purchase and the charging rise by the same amount, the demand constraints (1c) are unaffected; and because the stored level $\hat{h}_{r,1,\underline{t}_{n'},n'} = \tilde{h}_{r,1,\underline{t}_{n'},n'}$ is left unchanged, so is the remainder of the storage trajectory at n' . A negative residual is handled symmetrically by raising the discharge $\hat{h}_{r,1,\underline{t}_{n'},n'}^-$ above $\tilde{h}_{r,1,\underline{t}_{n'},n'}^-$, which disposes of the surplus carried-over energy and requires no purchase. Every remaining component of $\hat{y}$ coincides with that of $\tilde{y}$.

This correction preserves feasibility. No storage level is increased anywhere in the tree, so the storage-capacity constraints (1e) remain satisfied. Moreover, every corrective purchase lies in the first sub-period of a node: because each node in the final stage consists of more than one period, the first period $\underline{t}_{n'}$ of every node n' —including each leaf node—satisfies $\underline{t}_{n'} < T$. Since the emission constraints (1g) restrict exogenous purchases only in the final period T of the leaf nodes, these corrective purchases are unrestricted. The corrected solution $\hat{y}$ is therefore non-anticipative and feasible for (1).

The corrected solution $\hat{y}$ differs from the node-wise solution $\tilde{y}$ only in these additional purchases, which

offset the positive residuals. Restoring the balance at a child n' raises the purchase by $(\Delta_{r,n,n'})^+/\eta_r^+$, with the charging efficiency η_r^+ relating the purchased energy to the amount actually stored. Weighting each such purchase by its discounted unit cost $\pi_{n'} \beta^{(t_{n'}-1)} \lambda_{r,t_{n'}}$ and summing over all energy categories, non-leaf nodes, and their children bounds the objective increase of $\hat{y}$ over $\tilde{y}$ by (20). Because the objective of $\tilde{y}$ is no greater than that of y^* , the objective of $\hat{y}$ exceeds that of $(x_\omega^*, y_\omega^*)_{\omega \in \Omega}$ by at most (20), as claimed. $\square$

Rather than solving the node-wise problems and carrying out the construction described above, we recover non-anticipativity by solving a single LP. Because leaf nodes are not shared across scenario paths, the incumbent operational decisions y^* are already non-anticipative there, and only the internal nodes can exhibit violations. The LP fixes the investments at x^* and the leaf-node operational decisions at y^* —except those in each leaf’s first sub-period, which links the leaf to its parent node—and re-optimizes the remaining operational variables subject to the constraints of (1) and non-anticipativity. Because the non-anticipative solution constructed in Proposition 3 is feasible for this LP, the LP attains an objective that exceeds the incumbent’s by at most (20), a finite quantity since the storage installations—fixed at x^* and bounded in (1k)—bound the storage variables through the storage-capacity constraints (1e). In Section 5.2, we empirically observe that the correction step does not increase the objective value in any instance.

4.2. Out-of-Sample Performance Evaluation via Rolling-Horizon Simulation

The optimization model (1) is solved with nominal profiles for energy demand, renewable generation, and the heat pump coefficient of performance. In practice, these quantities may deviate from their nominal values. To assess how the resulting investment plans perform under such operational uncertainty, we conduct Monte Carlo simulations with realizations sampled from the distributions described in Section 3.2. Each realization is generated by adding either white- or pink-noise deviations to the nominal profiles, with pink noise producing temporally correlated deviations (Guevara et al., 2024). Given an investment plan, the dispatch decisions within a single replication are governed by the following problem:

$$\text{minimize } \left\{ \sum_{n \in \mathcal{N} \setminus \{0\}} \sum_{t \in \mathcal{T}_n} \sum_{r \in \mathcal{R}} \sum_{q \in \mathcal{Q}} \pi_n \beta^{(t-1)} \lambda_{r,t} p_{r,q,t,n} : (1b), (1c), (1d), (1e), (1f), (1l), (1m) \right\}. \quad (23)$$

Here, parameters $\Gamma_{j,r,q,t,t',n}$, $\mathbb{1}_{a,r,q,t,t',n}$, and $\delta_{r,q,t'}$ in constraints (1b) and (1c) are set to their sampled values. A complete simulation comprises many such replications, each under an independently sampled realization.

Solving problem (23) over the full horizon at once would assume perfect foresight of every future realization, whereas operational decisions are in fact made sequentially as uncertainty is revealed. To respect this information structure, we evaluate each replication through a sequence of LP dispatch problems solved over rolling windows that together span the planning horizon. Each window covers $\hat{l} \in \mathbb{Z}_+$ consecutive sub-periods: the first $\bar{l} \in \mathbb{Z}_+$ of these, with $\bar{l} \leq \hat{l}$, take their sampled realizations, while the remaining sub-periods retain their nominal values and serve as a planning lookahead. We solve the window’s dispatch problem, record the decisions of its first sub-period, and advance the window by one sub-period. The procedure repeats until the window has traversed every node and sub-period of the scenario tree.

To break ties among alternative optima, we assign a small negative objective coefficient to the stored-energy variables $h_{r,q,t,n}$ in the final sub-period of each window; this rewards carrying forward as much stored energy as possible without altering the primary objective. Emission constraints are omitted from the dispatch problems; any resulting net-zero shortfall is instead recorded and reported as the percentage of final-period demand met by exogenous purchases.

In the dispatch problems described so far, the lookahead sub-periods take nominal values. To make the dispatch decisions themselves robust to unfavorable conditions over the lookahead, we additionally construct a robust dispatch problem that mirrors the robust reformulation of Section 2.2. Within each window, the first $\bar{l}$ sub-periods retain their sampled realizations, while the lookahead sub-periods $(\bar{l} + 1, \dots, \hat{l})$ are displaced to the adverse end of their uncertainty intervals at robustness level ϵ . This provides a dispatch-side hedge that complements the robustness already embedded in the investment plans. In the Monte Carlo simulations, we evaluate both the nominal and the robust investment plans, dispatching each under both the nominal and the robust dispatch problems.

5. Computational Study

This section details the computational study of the proposed framework. Section 5.1 describes the campus-scale instance and the computational environment, Section 5.2 evaluates the computational performance against the benchmark methods, and Section 5.3 presents the case study for a university campus.

5.1. Experimental Setup

All computational experiments are conducted on a 64-bit Apple Silicon MacBook Pro running macOS Tahoe 26.2, equipped with an Apple M3 Pro processor (6 performance and 6 efficiency cores) and 36 GB of unified memory. The algorithms are implemented in Python 3.13.6, and all optimization problems are solved using Gurobi Optimizer 13.0.2 via its Python interface. The source code developed for this study is publicly available online (Şener et al., 2026).

The experiments are based on a clean electricity–heat transition planning instance constructed for the METU campus. In accordance with METU’s net-zero target for 2040, the study spans the 2026–2040 horizon. Exogenous energy procurement is therefore prohibited in the final period by setting $\gamma_r = 0$ for $r \in \mathcal{R}$. The planning horizon of $T = 15$ years is divided equally into $|\Psi| = 3$ stages of five years (periods), and each year is partitioned into sub-periods that aggregate two hours, yielding $|\mathcal{Q}| = 4,368$ sub-periods per year. An annual budget of \$20 million (M) is imposed in all periods.

Figure 2 presents a schematic of the integrated electricity–heat system formed by the candidate technologies considered for the transition. Photovoltaic panels and wind turbines are the electricity generation technologies, whereas parabolic troughs provide heat generation. Heat pumps serve as heat-transfer technologies that deliver heat using electricity, and their conversion efficiency is characterized by a coefficient of performance defined for each sub-period. Energy is stored by lithium-ion batteries on the electricity side and by hot water tanks on the heat side. Among these $|\mathcal{U}| = 6$ technologies, photovoltaic panels and lithium-ion batteries are subject to stochastic technological advancement, each with two advancement clusters, while the remaining four are treated as either mature or deterministically progressing. Two versions of photovoltaic panels with

different capacities are considered, both following the same technological progress. Table A.6 details the parameters of these candidate technologies.




Figure 2: Schematic of the integrated electricity-heat system.

The resulting scenario tree comprises $|\mathcal{N}| = 22$ nodes and $|\Omega| = 16$ scenario paths. In total, the corresponding optimization problem contains 9,435,001 constraints and 10,352,904 decision variables, of which 742 are investment variables and 318 are integer. The complete dataset used in the experiments is detailed in Appendix A.

5.2. Computational Performance

We evaluate the computational performance of the solution methodology proposed in Section 4.1 on the instance described in Section 5.1 at four operational temporal resolutions, defined by sub-periods that aggregate 24, 8, 4, and 2 hours. Each instance is solved to a relative optimality gap of 1%, subject to a time limit of 4 hours.

5.2.1. Linear Programming Relaxation

To demonstrate the effect of our algorithmic design choices on solution time, we first solve the LP relaxation of model (1) with different methods. Table 2 summarizes the performance of four approaches applied to the LP relaxation of model (1) across the considered temporal resolutions. For each method, it reports the best lower and upper bounds (LB and UB) on the objective value in millions of dollars, together with the resulting relative optimality gap (Gap) as a percentage. Solution times are given in seconds: a single total time (Time) for the *Extensive Form*, and the cumulative master-problem (Mstr.), subproblem (Sub.), and total (Total) times for each decomposition method. The four methods are defined as follows: *Extensive Form* solves the LP relaxation directly. *Algorithm 1* is the full decomposition algorithm, including both the valid inequalities (VIs) and the two-phase cut-addition strategy. *Algorithm 1 without VIs* differs from *Algorithm 1* only in that line 12 is removed, so that no valid inequalities are added to the master problem. *Single-Phase Algorithm 1* is obtained from *Algorithm 1* by removing lines 26–27, so that feasibility and optimality cuts are added together at every iteration without a phase transition. When Algorithm 1 and its variants are applied

to the LP relaxation, the integrality constraints are not reintroduced at the transition to Phase 2, so the two phases differ only in their cut-addition strategy.

Table 2: Comparison of solution methods for the LP relaxation of model (1) across temporal resolutions

	Extensive Form				Single-Phase Algorithm 1						Algorithm 1 without VIs						Algorithm 1					
	Time	LB	UB	Gap	Mstr.	Sub.	Total	LB	UB	Gap	Mstr.	Sub.	Total	LB	UB	Gap	Mstr.	Sub.	Total	LB	UB	Gap
24-h	107	122.53	122.53	0.00	2	38	41	122.51	122.53	0.02	2	39	41	122.51	122.53	0.02	2	31	34	122.51	122.56	0.04
8-h	3,450	155.79	155.79	0.00	17	411	432	155.77	155.81	0.03	8	312	323	155.33	156.20	0.56	8	281	292	155.33	156.01	0.43
4-h	14,400 [†]	–	–	–	57	1,588	1,656	157.69	157.69	0.00	11	911	927	157.06	158.18	0.71	12	882	901	156.93	158.08	0.73
2-h	14,400 [†]	–	–	–	178	8,084	8,330	161.12	161.13	0.00	17	3,290	3,320	160.27	161.82	0.96	23	3,710	3,757	160.32	161.77	0.89

[†]Time limit of 14,400 seconds (4 hours) reached. Dashes (–) indicate that no bound was available at termination.

Table 2 shows that, on the LP relaxation, the *Extensive Form* solves only the 24-hour and 8-hour instances to optimality, with solution times markedly longer than those of every decomposition-based variant. On the 4-hour and 2-hour instances, it reaches the time limit without returning any bound. All three variants of Algorithm 1, by contrast, solve every instance within the time limit. The proposed decomposition thus substantially outperforms the *Extensive Form*, particularly at the finer resolutions.

Although *Single-Phase Algorithm 1* employs the same 1% optimality gap as a stopping criterion, it consistently terminates with much smaller gaps, stopping as soon as it obtains its first incumbent solution. This behavior reflects how upper bounds arise in Benders decomposition: a valid upper bound becomes available only when all scenario-path subproblems are feasible for a given master solution. Because the cuts generated by different subproblems share common master variables, adding feasibility cuts alone would steer the master problem toward solutions that render all subproblems feasible. In the single-phase variant, however, optimality cuts are added at every iteration alongside the feasibility cuts, and they also steer the master solution toward minimizing the objective approximation. The master solution is therefore perturbed more strongly and cannot concentrate on producing solutions that are feasible across all subproblems.

Since iterations at which all subproblems are feasible occur only rarely, the lower bound continues to tighten across the intervening iterations. By the time such an iteration is reached, the lower bound has already converged close to the optimal value, leaving a near-zero gap. In Phase 2 of the two-phase strategy, by contrast, only feasibility cuts are added, so the master problem is guided toward solutions that are feasible across all subproblems, and incumbents are found more frequently. The single-phase variant therefore attains tighter final gaps than the two-phase variants, but at the cost of substantially longer solution times. This underscores the value of the two-phase strategy: Phase 1 promotes rapid lower-bound convergence by adding feasibility and optimality cuts together whereas Phase 2 focuses exclusively on feasibility cuts to find more incumbents.

The contribution of the valid inequalities can be assessed by comparing *Algorithm 1* against *Algorithm 1 without VIs*. Adding the valid inequalities reduces the total solution time at three of the four temporal resolutions, while the 2-hour instance exhibits the opposite behavior. On that instance, Phase 1 runs for 2,770 seconds and exits with a lower bound of \$159.6M when the valid inequalities are enabled, against 2,227 seconds and a weaker lower bound of \$156.9M when they are excluded.

Although the valid inequalities accelerate lower-bound convergence, enabling them on the 2-hour instance

delays the phase transition, causing the algorithm to spend more time in Phase 1. Once in Phase 2, the algorithm terminates in 978 seconds when the valid inequalities are enabled, compared with 1,088 seconds when they are disabled, consistent with the benefit of entering Phase 2 with a tighter lower bound. The longer Phase 1 nevertheless dominates, producing a higher total time. This outcome highlights the pivotal role of the phase-transition criterion (line 26) in determining the overall solution time: when the transition occurs later, the additional time spent in Phase 1 can outweigh the subsequent savings in Phase 2.

5.2.2. Mixed-Integer Linear Programming Model

Table 3 compares the solution methods for the original MILP model (1). These methods mirror those used for the LP relaxation, except that *Classical Benders* replaces the *Single-Phase* variant. *Classical Benders* applies Algorithm 1 with the same master–subproblem partitioning of the decision variables as in (15) and (17), assigning the investment variables to the master problem and the operational variables to the subproblem. Unlike the methodology proposed in Section 4, it does not apply the partial non-anticipativity relaxation, so the operational variables remain coupled across scenario paths and form a single subproblem rather than the scenario-wise subproblems of the proposed decomposition. The remaining three methods, together with the reported columns and their units, retain the definitions introduced for the LP relaxation.

Table 3: Comparison of solution methods for model (1) across temporal resolutions

	Extensive Form				Classical Benders						Algorithm 1 without VIs						Algorithm 1					
	Time	LB	UB	Gap	Mstr.	Sub.	Total	LB	UB	Gap	Mstr.	Sub.	Total	LB	UB	Gap	Mstr.	Sub.	Total	LB	UB	Gap
24-h	1,063	122.65	123.76	0.90	1,912	1,897	3,821	122.10	123.33	1.00	41	44	86	122.56	123.12	0.45	73	40	115	122.56	123.15	0.48
8-h	14,400 [†]	155.97	158.89	1.84	2,490	11,878	14,400 [†]	151.85	159.30	4.68	367	394	767	155.46	156.78	0.84	266	376	650	155.65	156.73	0.69
4-h	14,400 [†]	–	310.57	–	914	13,476	14,400 [†]	128.08	167.33	23.46	385	1,256	1,654	157.36	158.94	1.00	357	1,255	1,628	157.30	158.81	0.95
2-h	14,400 [†]	–	312.07	–	144	14,254	14,400 [†]	113.74	–	–	1,842	4,813	6,699	160.67	162.25	0.97	1,144	4,760	5,956	160.67	162.27	0.99

[†]Time limit of 14,400 seconds (4 hours) reached. Dashes (–) indicate that no bound was available at termination.

Table 3 shows that, apart from the 24-hour instance, neither the *Extensive Form* nor *Classical Benders* reaches the 1% optimality gap within the time limit. The *Extensive Form* solves only the 24-hour instance within this gap and returns no lower bound at the 4-hour and 2-hour resolutions. *Classical Benders* likewise solves only the 24-hour instance, where its solution time of 3,821 seconds exceeds even that of the *Extensive Form*. At the finer resolutions, it reaches the time limit with gaps of 4.68% or worse. Both variants of *Algorithm 1*, by contrast, solve every instance within the 1% optimality gap. On the 24-hour instance—the only one solved by either the *Extensive Form* or *Classical Benders*—they are roughly an order of magnitude faster than the *Extensive Form* and more than thirty times faster than *Classical Benders*.

Consistent with the LP-relaxation results, comparing *Algorithm 1* against *Algorithm 1 without VIs* shows that the valid inequalities reduce the total solution time at the three finer resolutions. On the 2-hour instance, identifying the most violated valid inequalities and separating them with Algorithm 2 takes 6.2 seconds over all iterations. At this finest resolution, *Algorithm 1* terminates in 1,330 iterations, whereas *Algorithm 1 without VIs* requires 1,460. Given this negligible separation cost, the valid inequalities accelerate convergence and meaningfully reduce the overall solution time, confirming their contribution to the performance of the solution procedure. We therefore continue with *Algorithm 1* in the remainder of the experiments.

When *Algorithm 1* is applied to the 2-hour instance, the LP that corrects the non-anticipativity violations of the operational decision variables (line 29) takes 4.7 seconds to solve. This correction step does not increase the objective value in any instance solved by the non-anticipativity-relaxed methods, for either the LP or the MILP model. Although Proposition 3 only provides an upper bound on this increment, in our experiments it is always zero. Overall, the proposed relaxation–correction strategy allows us to solve the original problem efficiently: relaxing non-anticipativity yields the decomposition that underlies the runtimes in Table 3, while restoring it afterward incurs only the negligible cost of the correction step.

5.3. Case Study

We now present the clean electricity–heat transition plan for the METU campus. We first describe the base-case investment plan, then evaluate the out-of-sample performance of the nominal and robust plans, and finally report sensitivity analyses on key problem parameters.

5.3.1. Base Case

Figure 3 presents the optimal investment plan obtained with *Algorithm 1* on the instance with 2-hour operational resolution described in Section 5.1. Each node of the scenario tree reports the capacity installed for each technology in every year, and the root node holds the first-stage decisions that are common to all scenario paths. Each branch label indicates the pace of technological progress, *s* for slow and *f* for fast, with the first letter referring to photovoltaic panels and the second to lithium-ion batteries. For instance, scenario path S_3 , denoted by $ss \times fs$, corresponds to a trajectory in which photovoltaic panels undergo slow progress followed by fast progress, while lithium-ion batteries undergo two consecutive stages of slow advancement.

The expected cost of the transition plan is \$162.27M, of which \$11.45M corresponds to electricity purchasing costs and \$23.64M to heat purchasing costs. The remaining \$127.18M consists of technology installation and O&M costs. In the first stage, the plan installs 12.0 MW of solar, 12.0 MW of wind, 1.5 MW h of battery storage, and 4.4 MW h of heat-pump capacity, all in the first year. Renewable generation and storage capacities are then expanded substantially in the second and third stages, driven both by the progress in technology costs and efficiencies and by the net-zero emission target enforced in the final period. Across the planning horizon, photovoltaic panels and lithium-ion batteries account for most of the installed capacity, whereas wind capacity is expanded only sporadically. Photovoltaic panels are preferred over wind turbines because of the limited wind potential at the installation site, where wind turbines achieve an annual capacity factor of only 10.4%, compared with 18.3% for photovoltaic panels. Battery storage grows sharply toward the end of the horizon as renewable penetration increases, with annual installations of up to 539.2 MW h. Once grid procurement is prohibited in the final period, substantial storage is needed to balance the intermittent output of renewable generation with demand.

Heat demand, on the other hand, is met mostly through electrified heating. Heat-pump installations appear in every stage and increase markedly toward the end of the planning horizon across all scenario paths. Parabolic troughs and hot water tanks, by contrast, are deployed only in small quantities and at a limited number of nodes, indicating that thermal energy generation and thermal storage play a marginal role relative to heat pumps powered by clean electricity. This observation is consistent with parabolic troughs requiring

further cost reductions to become economically competitive for widespread deployment (Khan et al., 2024).

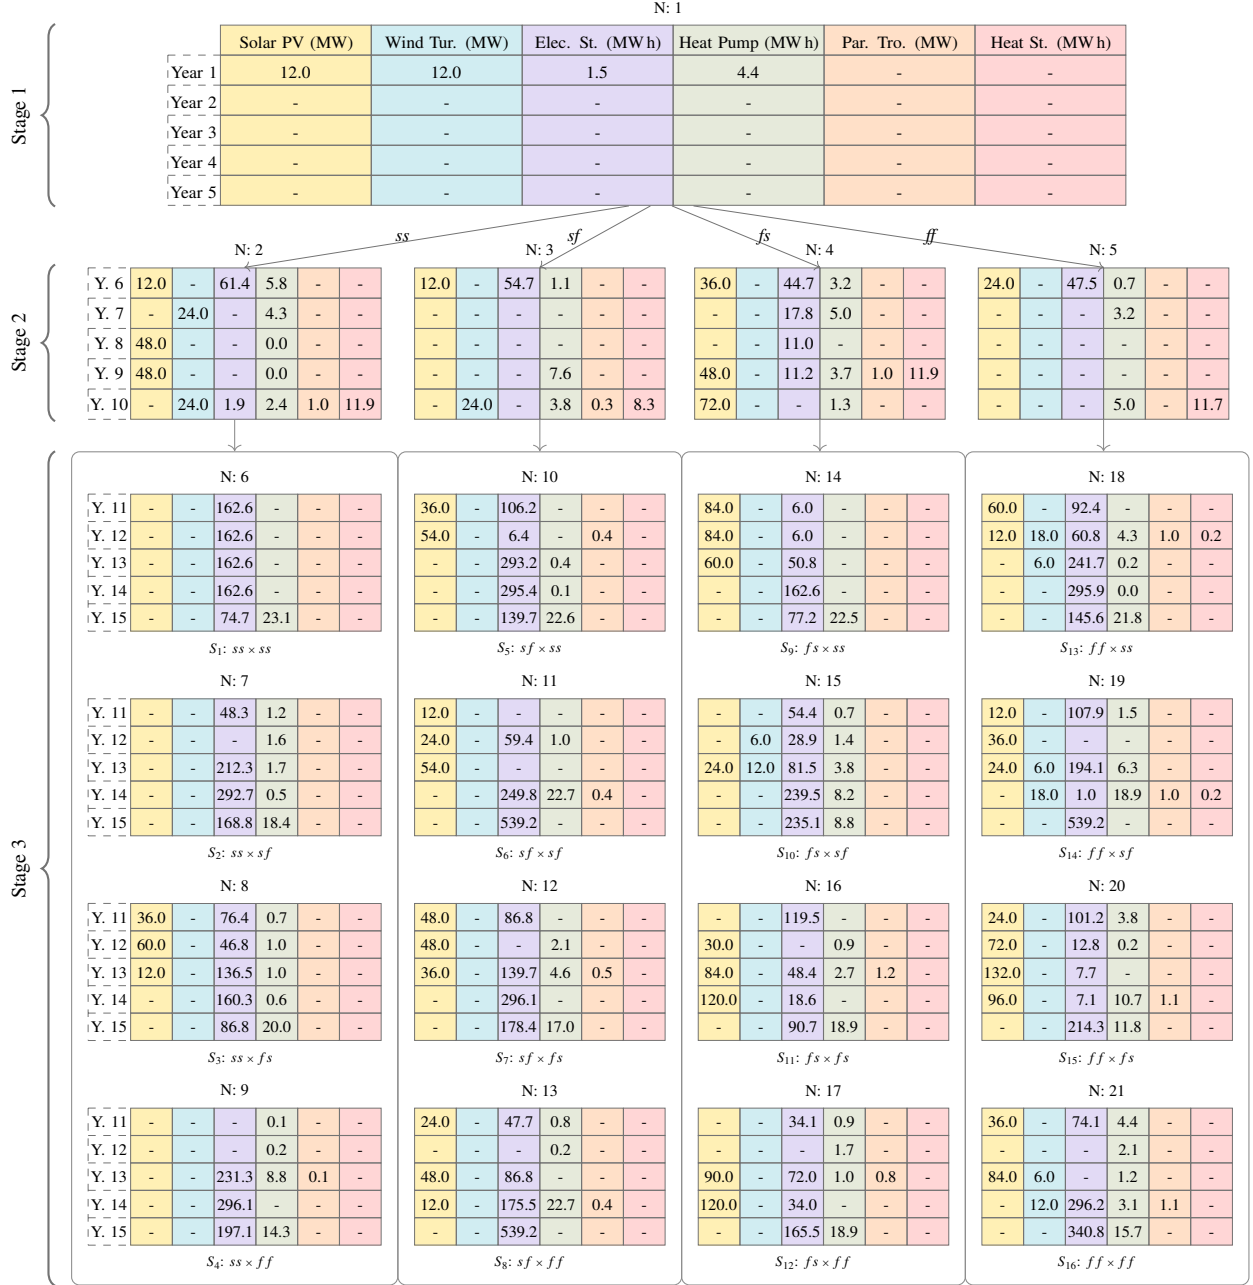


Figure 3: Decision tree for the base model.

The investment plans adapt to the realized technological advancements, differing across scenario paths in both the timing and the mix of installed technologies. The differences are clearest in solar deployment. Node 4, where solar advances quickly but lithium-ion batteries do not, installs 156 MW of solar already in the second stage, compared with 108 MW at node 2, where both technologies advance slowly; the four scenario paths descending from node 4 then expand photovoltaic capacity further in the final stage, reaching

the largest cumulative solar capacity in the tree, 402 MW on path S_{11} . Conversely, the slower branches invest less in solar: paths S_1 , S_2 , and S_4 , the children of node 2, add no photovoltaic capacity beyond the second stage and instead rely on storage, whose build-up is concentrated in the final stage and thereby shifts capacity toward the later years. The technology mix shifts as well; for example, wind turbine capacities are larger and installed earlier along trajectories with slower photovoltaic and lithium-ion progress.

These differences are also evident at the node level. Among the children of node 2 (nodes 6, 7, 8, and 9), only one node installs electricity generation technologies, whereas all other leaf nodes do. Budget utilization varies in a similar way: the budget constraints are active in every year of nodes 2, 6, 10, 14, and 18, a pattern most pronounced along the scenario paths in which technology development proceeds slowly during the transition from the second to the third stage. Along the first scenario path (S_1), for instance, the budget constraints remain active throughout the final decade, whereas no technology is installed in several years at nodes 3 and 5. Even leaf nodes branching from the same parent can exhibit considerably different technology mixes: given slow lithium-ion progress from the second to the third stage, photovoltaic installation capacities are 6 to 252 MW higher when the solar advancement in the second stage is fast. Together, these adaptations illustrate the sensitivity of the investment plan to the technological advancement trajectories and demonstrate the value of the multi-stage stochastic formulation, which tailors the timing and the technology mix to the evolving costs and efficiencies.

5.3.2. Out-of-Sample Performance Evaluation

Table 4 summarizes the out-of-sample performance of the nominal and robust investment plans, obtained by solving model (1) at investment robustness levels $\epsilon_i \in \{0, 0.1, 0.2\}$. Each plan is evaluated on a broader range of randomly generated operational data through the rolling-horizon Monte Carlo simulations described in Section 4.2. The operational deviations follow either a white-noise or a pink-noise process, and each plan is dispatched at robustness levels $\epsilon_d \in \{0, 1, 2\}$ applied to the lookahead sub-periods. In both the robust reformulation and the robust dispatch problems, the box uncertainty intervals are defined in terms of standard deviations: the deviation $\hat{v}_{d,i}$ that sets the half-width of each interval equals the standard deviation $s_{ds_{c,i}}$ derived in Section 3.2, so each uncertain parameter is displaced from its nominal value by ϵ standard deviations. The simulations use dispatch windows of $\hat{t} = 12$ sub-periods, of which the first $\bar{t} \in \{1, 3\}$ are treated as known.

Each configuration is evaluated over 100 replications solved in parallel. A single replication comprises 458,640 linear programs and takes approximately 550 seconds. The table reports the cost and quantity of purchased electricity and heat over the planning horizon, together with the net-zero shortfalls, expressed as the percentage of final-period electricity and heat demand met by exogenous purchases. All reported quantities are averaged across scenario paths and replications.

Increasing the investment robustness level ϵ_i trades a higher planned cost for more reliable out-of-sample performance. Building the plan under more conservative operational profiles raises its expected cost from \$162.27M for the nominal model ($\epsilon_i = 0$) to \$167.22M and \$172.81M for the robust models ($\epsilon_i = 0.1$ and $\epsilon_i = 0.2$), increases of 3.0% and 6.5%, respectively. In return, the out-of-sample purchasing cost falls by

roughly 8 to 9% in every configuration. The net-zero shortfalls respond in the same way: raising ϵ_i lowers them in nearly all configurations, the only exception being the already negligible electricity shortfall under white noise with $\bar{l} = 3$, where the heat shortfall nonetheless continues to decline. These shortfalls remain small in absolute terms throughout, so the plans rarely miss the net-zero target by a wide margin. The gains are uneven across commodities, however: heat shortfalls remain consistently larger than electricity shortfalls and are harder to drive to zero. This mirrors the smaller proportional reduction in heat purchases and identifies heat as the binding commodity for the installed technology mix.

Table 4: Out-of-sample Monte Carlo simulation results for the nominal and robust investment plans under white- and pink-noise deviations.

Investment ϵ_i	Noise	Dispatch ϵ_d	$\bar{l} = 1$					$\bar{l} = 3$				
			Purchasing cost (\$M)	Purchased (MWh)		Shortfall (%)		Purchasing cost (\$M)	Purchased (MWh)		Shortfall (%)	
				Electricity	Heat	Electricity	Heat		Electricity	Heat	Electricity	Heat
0	White	0	35.03	88,948	746,917	0.152	0.398	34.58	87,448	736,730	0.034	0.072
		1	34.76	86,775	745,658	0.006	0.064	34.56	86,759	738,620	0.002	0.039
		2	35.08	86,497	757,951	0.000	0.078	34.70	86,556	744,526	0.000	0.046
	Pink	0	35.78	93,550	757,007	1.309	2.363	35.00	89,784	742,430	0.251	0.839
		1	35.16	88,868	751,813	0.202	0.877	34.88	88,420	743,380	0.049	0.613
		2	35.38	88,083	762,360	0.037	0.637	35.01	88,113	749,050	0.027	0.592
0.1	White	0	33.08	78,120	728,475	0.131	0.281	32.55	76,166	717,145	0.037	0.049
		1	32.76	75,525	727,145	0.004	0.033	32.51	75,447	718,704	0.002	0.017
		2	33.10	75,259	740,205	0.000	0.035	32.65	75,293	724,433	0.000	0.019
	Pink	0	33.87	83,172	737,764	1.088	1.967	33.01	78,928	722,733	0.242	0.648
		1	33.19	77,966	732,764	0.153	0.641	32.87	77,435	723,175	0.039	0.426
		2	33.43	77,154	744,097	0.025	0.433	32.99	77,157	728,608	0.019	0.405
0.2	White	0	31.93	73,040	709,055	0.121	0.244	31.43	71,443	697,862	0.038	0.047
		1	31.65	70,664	708,453	0.004	0.022	31.39	70,667	699,557	0.002	0.010
		2	32.04	70,314	723,484	0.000	0.017	31.55	70,395	706,128	0.000	0.010
	Pink	0	32.73	77,890	719,404	1.034	1.767	31.91	74,125	704,234	0.237	0.563
		1	32.10	73,037	714,874	0.143	0.544	31.77	72,651	704,730	0.034	0.349
		2	32.39	72,200	728,013	0.021	0.348	31.90	72,256	710,971	0.015	0.325

Beyond the choice of investment plan, the temporal structure of the operational deviations has a pronounced effect on out-of-sample performance. Under white noise, the deviations are temporally independent and largely offset one another across sub-periods whereas under pink noise, unfavorable realizations, such as below-nominal renewable generation and above-nominal demand, can persist over consecutive sub-periods, leaving the installed technology mix unable to fully meet demand. Pink noise therefore raises purchased electricity and heat in every configuration. Two levers mitigate this effect: Extending the known portion of the window from $\bar{l} = 1$ to $\bar{l} = 3$ lets each dispatch problem anticipate its upcoming sub-periods. When no dispatch robustness is applied, this more than halves the shortfalls in every configuration and also lowers the purchased quantities, especially under pink noise. Dispatch-side robustness provides a complementary hedge that leaves the investment plan unchanged, although its effect on both the net-zero shortfalls and the purchasing cost is non-monotone. The first increment ($\epsilon_d = 1$) reduces both the shortfalls and the purchasing cost in every configuration, whereas the second ($\epsilon_d = 2$) dispatches so conservatively that it raises heat purchases enough to increase the overall purchasing cost. This reversal arises because the second increment doubles the width

of the dispatch problem’s uncertainty intervals, bracing each lookahead sub-period against more pessimistic combination of low renewable generation and a low heat-pump coefficient of performance. Anticipating these harsh conditions, the dispatch draws down its stored electricity in the current sub-periods rather than holding it in reserve, leaving less available for the periods that follow. When these conservative expectations go unrealized, the depleted storage can no longer drive the heat pump, so the unmet heat demand is covered by exogenous purchases, which raises the overall purchasing cost.

5.3.3. Sensitivity Analyses

We analyze the effect of three key problem parameters on the transition plan: the final-period emission target, the investment budget, and the energy demand.

The Effect of Emission Target: The first analysis concerns the sensitivity of the expected cost of the transition plan to the final-period emission target. The base-case instance is re-solved with $\gamma_r \in \{0.01, 0.02, 0.03, 0.04, 0.05\}$ for all $r \in \mathcal{R}$, allowing the corresponding fraction of final-period demand in each energy category to be met by exogenous purchases. The validity of the inequalities (3) relies on the elimination of the exogenous procurement decision variables, which is possible only when $\gamma_r = 0$. These instances are therefore solved using *Algorithm 1 without VIs*.

Table 5 reports the expected cost across the range of exogenous procurement allowances. The results reveal a clear trade-off between the stringency of the emission target and the cost of the transition plan: permitting exogenous purchases for just 1% of final-period demand lowers the expected cost from the base-case \$162.27M to \$125.67M, while each additional percentage point yields a progressively smaller reduction, with the cost reaching \$108.80M at the 5% allowance. The large saving from the first percentage point indicates that insisting on an exact net-zero target carries a substantial cost premium. The allowances also reshape the technology mix, with the initial allowances substantially reducing the installed storage capacity and further increments primarily lowering the installed capacities of generation technologies and heat pumps. This occurs since purchases meet a larger share of demand in sub-periods with unfavorable renewable performance, lowering the capacity that would otherwise be required to serve it.

Table 5: Expected costs under different final-period exogenous procurement allowances.

$\gamma_r, r \in \mathcal{R}$	0	0.01	0.02	0.03	0.04	0.05
Expected cost (\$M)	162.27	125.67	118.61	114.51	111.36	108.80

The Effect of Investment Budget: The annual budget, fixed at \$20M in every period in the base case, is varied next. Raising it to \$25M reduces the expected cost from the base case to \$153.38M, and a further increase to \$30M lowers it to \$149.07M. The additional reduction becomes smaller as the budget grows. These diminishing marginal savings reflect the budget acting as a binding constraint: as noted for the base-case decision tree (Figure 3), the budget caps are active in many years, particularly along the scenario paths with slow technology development. Relaxing the annual budget loosens these caps, so that capacity can be installed closer to when it is needed and at lower cost, rather than earlier to remain within the annual limit.

The Effect of Energy Demand: Although no clear driver of demand growth is anticipated for the METU

campus, the robustness of the transition plan to such growth is nonetheless examined by increasing the nominal electricity and heat demand at compounding annual rates of 0.5% and 1% over the planning horizon. The expected cost rises from the base case to \$180.32M under 0.5% growth and to \$202.89M under 1% growth, with the increase accelerating at the higher growth rate. Because the demand increments compound over the horizon and exogenous purchases are prohibited in the final period under the net-zero target, the additional final-period demand must be met entirely by the installed capacity. The plan therefore installs considerably more renewable generation and storage capacity. Since the budget binds during the final decade along some scenario paths, part of this additional capacity is installed in the early years, when technologies are more expensive, which adds further to the expected cost.

6. Conclusion

This paper studies clean electricity–heat transition planning for a campus-scale integrated energy system, where long-term investment decisions must adapt to uncertain technological progress while remaining compatible with high-resolution operations. We represent this setting with a multi-stage stochastic mixed-integer program and develop a robust reformulation that captures operational uncertainty through box uncertainty sets. To solve the resulting large-scale instances, we propose a Benders decomposition algorithm with partial non-anticipativity relaxation: operational non-anticipativity is relaxed during the decomposition and restored only at termination. The algorithm is further strengthened with valid inequalities and a two-phase cut-addition strategy.

The METU campus case study shows that the proposed algorithm substantially improves computational efficiency. It solves every instance, across all temporal resolutions, to a 1% optimality gap within the time limit, whereas both the extensive form and classical Benders decomposition solve only the coarsest 24-hour instance. On that instance, the proposed algorithm is roughly an order of magnitude faster than the extensive form and more than thirty times faster than classical Benders. The resulting net-zero transition plan has an expected cost of \$162.27M and relies primarily on photovoltaic panels, lithium-ion batteries, and heat pumps.

The rolling-horizon Monte Carlo simulations and sensitivity analyses provide several planning insights. Increasing the investment robustness level raises the planned transition cost, but it reduces out-of-sample purchasing costs and net-zero shortfalls. Temporally correlated operational deviations are more challenging than independent deviations, especially for heat supply, while dispatch-side robustness can mitigate these effects when its conservatism is chosen carefully. The final-period emission target is the dominant cost driver: enforcing an exact net-zero target, rather than allowing even a marginal final-period shortfall, carries a substantial cost premium.

Beyond the campus application, the methodology applies to multi-stage stochastic mixed-integer programs with a similar master–subproblem structure. In such problems, the master problem retains variables whose non-anticipativity is difficult to restore, while the decomposition assigns the remaining variables to scenario-wise subproblems and corrects their non-anticipativity after termination. The correction step must recover a feasible non-anticipative solution while limiting any change in the objective value. In the present model, the correction LP satisfies this requirement and restores non-anticipativity without increasing the objective

value across the computational instances. A natural next step is to scale the framework to larger multi-energy networks, supporting transition planning beyond the campus setting.

Acknowledgments

This work was supported by the Scientific and Technological Research Council of Turkey [grant number 222M243].

References

- Azimian, M., Shen, X., & Gharehpetian, G. B. (2025). Robust scenario-based stochastic expansion planning of multi-carrier microgrids considering incentive-based loans. *Applied Energy*, 401, 126593. <https://doi.org/10.1016/j.apenergy.2025.126593>.
- Bacci, L., Dal Cin, E., Carraro, G., Rech, S., & Lazzaretto, A. (2025). Economic–energy–environmental optimization of a multi-energy system in a university district. *Energies*, 18, 413. <https://doi.org/10.3390/en18020413>.
- Bentley, J. (1984). Programming pearls: Algorithm design techniques. *Communications of the ACM*, 27, 865–873. <https://doi.org/10.1145/358234.381162>.
- Birge, J. R. (1985). Decomposition and partitioning methods for multistage stochastic linear programs. *Operations Research*, 33, 989–1007. <https://doi.org/10.1287/opre.33.5.989>.
- Birge, J. R., & Louveaux, F. (1997). *Introduction to stochastic programming*. New York: Springer.
- Chen, Y., Sergi, B., Ho, J., Stephen, G., Cole, W., & Powell, K. (2025). Sparse chronology strategy for integrating seasonal energy storage in capacity expansion models. *Journal of Energy Storage*, 132, 117950. <https://doi.org/10.1016/j.est.2025.117950>.
- Davis, S. J., Lewis, N. S., Shaner, M., Aggarwal, S., Arent, D., Azevedo, I. L., Benson, S. M., Bradley, T., Brouwer, J., Chiang, Y.-M. et al. (2018). Net-zero emissions energy systems. *Science*, 360, eaas9793. <https://doi.org/10.1126/science.aas9793>.
- Dupačová, J., Consigli, G., & Wallace, S. W. (2000). Scenarios for multistage stochastic programs. *Annals of Operations Research*, 100, 25–53. <https://doi.org/10.1023/A:1019206915174>.
- Dupačová, J., Gröwe-Kuska, N., & Römis, W. (2003). Scenario reduction in stochastic programming. *Mathematical Programming*, 95, 493–511. <https://doi.org/10.1007/s10107-002-0331-0>.
- Fang, Y., Han, J., Du, E., Jiang, H., Fang, Y., Zhang, N., & Kang, C. (2024). Electric energy system planning considering chronological renewable generation variability and uncertainty. *Applied Energy*, 373, 123961. <https://doi.org/10.1016/j.apenergy.2024.123961>.
- Füllner, C., Sun, X. A., & Rebennack, S. (2024). A new framework to generate Lagrangian cuts in multistage stochastic mixed-integer programming. Optimization Online. <https://optimization-online.org/2024/08/a-new-framework-to-generate-lagrangian-cuts-in-multistage-stochastic-mixed-integer-programming/>.
- Gao, N., Gao, D. W., & Fang, X. (2023). Manage real-time power imbalance with renewable energy: Fast generation dispatch or adaptive frequency regulation? *IEEE Transactions on Power Systems*, 38, 5278–5289. <https://doi.org/10.1109/TPWRS.2022.3232759>.
- Gries, D. (1982). A note on a standard strategy for developing loop invariants and loops. *Science of Computer Programming*, 2, 207–214. [https://doi.org/10.1016/0167-6423\(83\)90015-1](https://doi.org/10.1016/0167-6423(83)90015-1).
- Guevara, E., Babonneau, F., & Homem-de-Mello, T. (2024). Uncertainty dynamics in energy planning models: An autoregressive and Markov chain modeling approach. *Computers & Industrial Engineering*, 191, 110084. <https://doi.org/10.1016/j.cie.2024.110084>.
- Higle, J. L., & Sen, S. (1991). Stochastic decomposition: An algorithm for two-stage linear programs with recourse. *Mathematics of Operations Research*, 16, 650–669. <https://doi.org/10.1287/moor.16.3.650>.

- Huylo, M., Taheri, S., & Novoselac, A. (2025). Integration of renewable energy generation and storage systems for emissions reduction in an islanded campus microgrid. *Building and Environment*, 274, 112736. <https://doi.org/10.1016/j.buildenv.2025.112736>.
- Ioannou, A., Fuzuli, G., Brennan, F., Yudha, S. W., & Angus, A. (2019). Multi-stage stochastic optimization framework for power generation system planning integrating hybrid uncertainty modelling. *Energy Economics*, 80, 760–776. <https://doi.org/10.1016/j.eneco.2019.02.013>.
- IRENA (2025). Renewable power generation costs in 2024. International Renewable Energy Agency, Abu Dhabi. <https://www.irena.org/Publications/2025/Jul/Renewable-power-generation-costs-in-2024> (ISBN 978-92-9260-669-5).
- Khan, M. I., Gutiérrez-Alvarez, R., Asfand, F., Bicer, Y., Sgouridis, S., Al-Ghamdi, S. G., Jouhara, H., Asif, M., Kurniawan, T. A., Abid, M., Pesyridis, A., & Farooq, M. (2024). The economics of concentrating solar power (CSP): Assessing cost competitiveness and deployment potential. *Renewable and Sustainable Energy Reviews*, 200, 114551. <https://doi.org/10.1016/j.rser.2024.114551>.
- Kleywegt, A. J., Shapiro, A., & Homem-de-Mello, T. (2002). The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization*, 12, 479–502. <https://doi.org/10.1137/S1052623499363220>.
- Li, C., Conejo, A. J., Sirola, J. D., & Grossmann, I. E. (2022). On representative day selection for capacity expansion planning of power systems under extreme operating conditions. *International Journal of Electrical Power & Energy Systems*, 137, 107697. <https://doi.org/10.1016/j.ijepes.2021.107697>.
- Marcy, C., Goforth, T., Nock, D., & Brown, M. (2022). Comparison of temporal resolution selection approaches in energy systems models. *Energy*, 251, 123969. <https://doi.org/10.1016/j.energy.2022.123969>.
- Mavromatidis, G., & Petkov, I. (2021). MANGO: A novel optimization model for the long-term, multi-stage planning of decentralized multi-energy systems. *Applied Energy*, 288, 116585. <https://doi.org/10.1016/j.apenergy.2021.116585>.
- Moore, G. E. (1998). Cramming more components onto integrated circuits. *Proceedings of the IEEE*, 86, 82–85. <https://doi.org/10.1109/JPROC.1998.658762>.
- Mulvey, J. M., & Ruszczyński, A. (1995). A new scenario decomposition method for large-scale stochastic optimization. *Operations Research*, 43, 477–490. <https://doi.org/10.1287/opre.43.3.477>.
- NREL (2025). System Advisor Model (SAM), Version 2025.4.16. National Renewable Energy Laboratory, Golden, CO, USA. <https://sam.nrel.gov> (Accessed July 1, 2026).
- Pereira, M. V. F., & Pinto, L. M. V. G. (1991). Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming*, 52, 359–375. <https://doi.org/10.1007/BF01582895>.
- Rathi, T., & Zhang, Q. (2022). Capacity planning with uncertain endogenous technology learning. *Computers & Chemical Engineering*, 164, 107868. <https://doi.org/10.1016/j.compchemeng.2022.107868>.
- Rockafellar, R. T., & Wets, R. J.-B. (1991). Scenarios and policy aggregation in optimization under uncertainty. *Mathematics of Operations Research*, 16, 119–147. <https://doi.org/10.1287/moor.16.1.119>.
- Şener, A. E., Kocuk, B., & Yüksel, T. (2025). A dynamic strategic plan for transition to campus-scale clean electricity using multi-stage stochastic programming. *arXiv:2506.16423*.
- Şener, A. E., Kocuk, B., & Yüksel, T. (2026). stochastic-campus-energy-transition. GitHub repository. <https://github.com/emrseners/stochastic-campus-energy-transition> (Accessed July 1, 2026).
- Teichgraeber, H., & Brandt, A. R. (2019). Clustering methods to find representative periods for the optimization of energy systems: An initial framework and comparison. *Applied Energy*, 239, 1283–1293. <https://doi.org/10.1016/j.apenergy.2019.02.012>.
- Tian, X., Zhou, Y., Morris, B., & You, F. (2022). Sustainable design of Cornell University campus energy systems toward climate neutrality and 100% renewables. *Renewable and Sustainable Energy Reviews*, 161, 112383. <https://doi.org/10.1016/j.rser.2022.112383>.

Vergara-Zambrano, J., Brahmhatt, P., & Avraamidou, S. (2025). A multi-scale optimization framework for energy transition planning in urban areas: Insights from a university campus case study. *Sustainable Energy, Grids and Networks*, 44, 101996. <https://doi.org/10.1016/j.segan.2025.101996>.

Zakaria, A., Ismail, F. B., Lipu, M. S. H., & Hannan, M. A. (2020). Uncertainty models for stochastic optimization in renewable energy applications. *Renewable Energy*, 145, 1543–1571. <https://doi.org/10.1016/j.renene.2019.07.081>.

Zou, J., Ahmed, S., & Sun, X. A. (2019). Stochastic dual dynamic integer programming. *Mathematical Programming*, 175, 461–502. <https://doi.org/10.1007/s10107-018-1249-5>.

Appendix A. Data

Table A.6 presents the annual energy yields of the candidate renewable technologies, computed from hourly generation profiles simulated with the National Renewable Energy Laboratory’s System Advisor Model (SAM) (NREL, 2025). For each technology, SAM uses the geographic location of the METU campus and the technical specifications listed in Table A.6 to simulate hourly energy output over a representative meteorological year. The simulated generation profiles, demand series, and detailed input data used to construct the model are publicly available in the source code repository (Şener et al., 2026).

Table A.6: Candidate renewable technologies (costs are in 2024 \$).

Technology	Capacity	Installation Cost	Annual O&M Cost [†]	Footprint	Annual Energy Yield	Economic Lifetime	Annual Degradation Rate	Round-trip Efficiency	Self-discharge Rate	Reference
Solar PV ₁	6 MW	\$3,190,170	8.35 \$/kW – 0.918	28,370 m ²	9.63 GW h	25 years	0.5%	-	-	Şener et al. (2025)
Solar PV ₂	12 MW	\$5,832,190	13.74 \$/kW – 0.898	56,739 m ²	19.26 GW h	25 years	1.6%	-	-	Şener et al. (2025)
Wind Turbine ₁	6 MW	\$5,430,435	13.99 \$/kW	5,280 m ²	4.15 GW h	25 years	0.2%	-	-	NREL (2025)
Parabolic Trough ₁	4.02 MW	\$2,812,480	9.18 \$/kW h	12 m ²	-	20 years	2.0%	-	-	Mavromatidis & Petkov (2021)
Heat Pump ₁	246 kW h	\$150,563	5.50 \$/kW h	30.0 m ²	-	20 years	1.0%	90.0%	0.06%/h	Şener et al. (2025)
Lithium-ion Battery ₁	1 MW h	\$305,556	0.30 \$/kW h	8.7 m ²	-	30 years	-	90.0%	0.50%/h	Mavromatidis & Petkov (2021)
Hot Water Storage ₁	1 MW h	\$15,000								

[†]O&M cost shown as: initial annual cost (\$/kW) – annual reduction multiplier

Table A.7: Five-year technology advancement multipliers.

Technology	Advancement Probability	Cost Multiplier	Efficiency Multiplier	Reference
Solar photovoltaic	Fast (0.667) Slow (0.333)	0.551 0.856	1.133 1.056	Şener et al. (2025)
Lithium-ion battery	Fast (0.464) Slow (0.536)	0.348 0.634	1.000 1.000	Şener et al. (2025)
Wind turbine	Medium (1.000)	0.815	1.000	Şener et al. (2025)
Parabolic trough	Medium (1.000)	0.683	1.000	IRENA (2025)
Heat pump	Medium (1.000)	0.874	1.156 [†]	Mavromatidis & Petkov (2021)
Thermal storage	Medium (1.000)	1.000	1.000	Mavromatidis & Petkov (2021)

[†]The advancement multipliers are applied multiplicatively across successive five-year periods, whereas the heat pump efficiency multiplier is applied additively.