A Multistage Stochastic Optimization Framework for Long-Term District Heating Capacity Expansion

1st Ali Kök
Energy Economic Group
Technische Universität Wien
Vienna, Austria
koek@eeg.tuwien.ac.at

2nd Jonathan Hachez
Thermo and Fluid Dynamics (FLOW)
Vrije Universiteit Brussel
Brussels, Belgium
jonathan.hachez@vub.be

3rd Lukas Kranzl
Energy Economic Group
Technische Universität Wien
Vienna, Austria
kranzl@eeg.tuwien.ac.at

***Abstract*—District heating (DH) systems play a pivotal role in Europe's decarbonization agenda, especially amid new EU regulations mandating local heating and cooling plans for municipalities above 45,000 inhabitants, with a target to increase renewables and excess heat in DH by 2.1 ppt annually. However, uncertainties such as demand growth and fuel prices present major challenges for long-term planning. This paper develops a multistage stochastic capacity expansion model for long-term DH investments, using a Markov chain for annual demand evolution and stagewise-independent draws for fuel prices. The approach adopts Stochastic Dual Dynamic Programming (SDDP), implemented in Julia/SDDP.jl, separating investment and operational stages while accounting for lead times and salvage values. A hypothetical Central European case study dominated by gas CHP and boilers illustrates the model's capability. Results suggest a multistage formulation can significantly improve robustness and cost-effectiveness in DH planning under deep uncertainty.**

***Index Terms*—Stochastic Dual Dynamic Programming, District Heating, Capacity Expansion, Markov Chain**

I. Introduction

District heating (DH) systems are pivotal for decarbonizing the buildings and heating sector as they can leverage renewable energy and excess heat sources. Recent policies in EU now mandate that municipalities with more than 45,000 inhabitants develop local heating and cooling plans, with specific targets to increase the share of renewables and excess heat in DH networks by 2.1 percentage points annually [1]. This drives a growing need for cost-minimal pathways to expand and decarbonize DH systems, yet municipal decision-makers face considerable uncertainties in future demand, fuel and electricity prices, and policy support.

While DH modeling has become increasingly advanced, most studies that incorporate uncertainty focus on short-term operational or bidding aspects, with limited temporal scope. For example, [2] use Monte Carlo simulation to optimize day-ahead operation of cogeneration plants under uncertain electricity prices, and [3] apply particle swarm optimization in a district energy system considering stochastic renewable power generation. Similarly, [4] investigate the techno-economic performance of fifth-generation DH using Monte Carlo simulations, analyzing the levelized cost of heat under varying electricity prices and substation costs. These works, along with additional research on thermal network design and control [5]–[7], highlight how uncertainties can influence DH planning and operation in the short term.

Nevertheless, relatively few studies address long-term capacity expansions of DH supply under evolving demand and policy constraints. Existing research on DH expansion typically uses scenario-based deterministic models or two-stage frameworks that treat investment as a single decision point. Examples include [8], where day-ahead bidding under uncertain electricity prices is considered but not multiperiod investments, and [9], which analyse two-stage stochastic planning for multi-energy systems but do not incorporate demand evolution or salvage valuations. In contrast, recent work in the electric power and hydropower sectors—including hydro scheduling and mining expansions—demonstrates the effectiveness of Stochastic Dual Dynamic Programming (SDDP) for large-scale, multistage investments where uncertainty evolves over many years. The policy-graph method introduced by [10] formalizes a way to represent Markovian states and stagewise-independent processes, enabling tractable scenario expansions. Additionally, [11] discuss SDDP can be used to model capacity planning of renewable energy systems.

The work presented here applies the policy graph framework directly to DH infrastructure planning in a genuinely multi-stage framework using SDDP,jl [12]: expansions are chosen at discrete intervals (2020, 2030, 2040, 2050), each followed by an operational phase under uncertain energy prices and demand growth. We use a Markov chain to model annual demand evolution, stagewise-independent random fuel prices, and realistic lead times that delay new capacity by one model year. Any capacity extending beyond the final stage earns a salvage credit, avoiding the penalization of expansions that remain beneficial past the horizon.

The paper first describes the problem's key components and SDDP solution approach, then details how we handle demand, lead times, salvage, and price uncertainty. We conclude by applying this framework to hypothetical case study, and discuss how the results can guide robust long-term DH strategies.

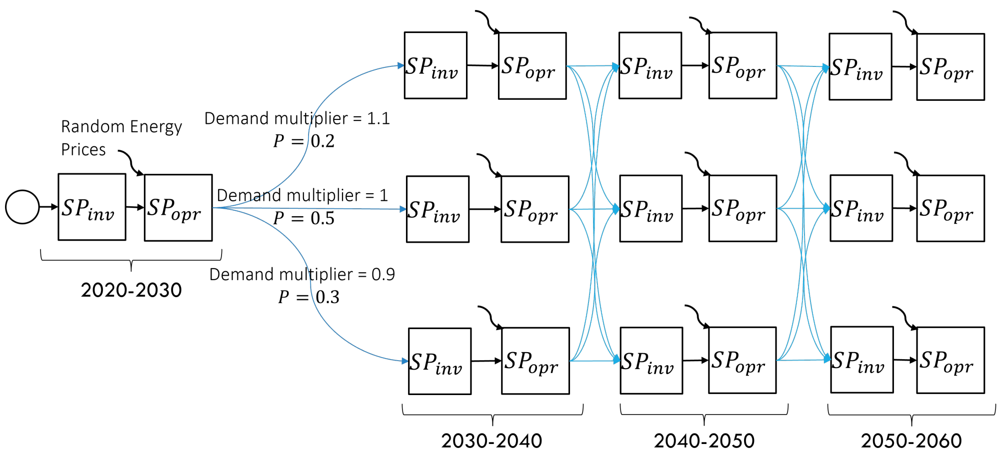

Fig. 1. Illustration of the policy graph structure with deterministic transitions (black edges) in the initial years and Markovian branches (blue edges) in subsequent stages. Wiggly lines represent stagewise-independent random energy prices.

II. Methodology

A. Policy Graph Representation

To efficiently model multistage decision-making under uncertainty, we use the policy graph framework, proposed by [10], a structured approach to defining state transitions and decision dependencies within the Stochastic Dual Dynamic Programming (SDDP) framework.

A policy graph is a directed acyclic graph where nodes represent subproblems in which an agent makes a decision, and edges denote (stochastic) state transitions between nodes. Probabilistic Branching reflects the stochastic nature of exogenous variables, such as energy prices and demand growth. A policy graph can be defined as a tuple $\mathcal{G} = (\mathcal{N}, \mathcal{E}, \Phi)$, where $\mathcal{N}$ is the set of nodes, representing decision stages, $\mathcal{E}$ is the set of directed edges, denoting feasible transitions, Φ_{ij} represents the transition probability from node i to node j [11]. Each node i has an associated subproblem:

$$SP_i(x, \omega) = \min_{u, x'} C_i(x, u, x', \omega), \tag{1}$$

subject to:

$$(x', u) \in \mathcal{X}_i(x, \omega), \tag{2}$$

where x represents the incoming states variable from previous node, x' is the outgoing state variables to the subsequent node, u denotes the control variables, ω is the stochastic realization of exogenous parameters, $\mathcal{X}_i(x, \omega)$ defines feasibility constraints [11]. The conventional cost-to-go functions are not directly represented; instead, they can be derived implicitly from the policy graph's structure.

We consider a horizon of T stages (an even number), where odd t are investment stages and even t are operational. Each pair $(t, t+1)$ forms a model year. We define $\mathcal{T}_{\text{inv}} = \{1, 3, \dots, T-1\}$, $\mathcal{T}_{\text{opr}} = \{2, 4, \dots, T\}$. A capacity expansion at stage $t \in \mathcal{T}_{\text{inv}}$ becomes available at $t+2$ (one-year lead). In each operational stage $t \in \mathcal{T}_{\text{opr}}$, we optimize hourly dispatch given the alive capacities.

Figure 1 shows the proposed policy graph structure, where the first model year has a deterministic annual demand while the subsequent model years follow a Markov chain with three possible states: an increase in demand, stable demand, or demand reduction. Φ_{ij} represents the probability of transitioning between these states. Demand evolution is modelled using demand multiplier with realisations ξ_t. Blue edges represent stochastic transitions between Markovian demand states.

Energy prices are modelled as stage-wise independent random variables and realised at each operational stage $t \in \mathcal{T}_{\text{opr}}$. We let ω_t be the realisation of the random energy price in stage t. At the investment stage, energy prices remain uncertain. At each investment stage, the decision-maker knows the current demand state and past investment choices but lacks precise knowledge of future annual demand realisations and energy prices.

B. Investment Subproblem

At each investment stage $t \in \mathcal{T}_{\text{inv}}$, the model determines how much new capacity to add for each technology j. We let $u^{\text{inv}}_{j,t} \geq 0$ denote the newly built capacity in stage t. Additionally, we track all vintage capacities $x^{\text{vintage}}_{j,v,t}$, which represent the capacity of technology j that was originally installed in stage v, where $v \in \{v \mid v \in \{0\} \cup \mathcal{T}_{\text{inv}}, v \leq t\}$. This definition ensures that vintage stages include the terminal stage 0, representing existing capacities, as well as all past investment stages up to and including the current stage t.

To capture lead times and lifetimes, we define the alive capacity:

$$x^{\text{alive}}_{j,t} = \sum_{v=0}^{t} x^{\text{vintage}}_{j,v,t} \cdot \lambda_{j,v,t}, \tag{3}$$

where $\lambda_{j,v,t}$ is a binary parameter indicating whether technology j built in stage v is operational at stage t, defined as

$$\lambda_{j,v,t} = \begin{cases} 1, & \text{if } \lceil \frac{t}{2} \rceil - \lceil \frac{v}{2} \rceil \geq 1, \lceil \frac{t}{2} \rceil - \lceil \frac{v}{2} \rceil \leq \left\lceil \frac{L_j}{2} \right\rceil, \\ 0, & \text{otherwise,} \end{cases} \tag{4}$$

where L_j is the the operational lifetime of technology j, in terms of number of stages. The ceiling operators are used to convert the number of stages into model years, as a pair of stages represents a model year. The first condition in (4) ensures that the investment is realized after one model year while the second condition ensures that the technology is within its operational lifetime.

To track the annual demand evolution under the Markov chain, we introduce the state variable x_t^d that stores the value of the annual demand multiplier.

To account for the residual value of capacities that remain operational beyond the planning horizon, we include a salvage value term:

$$\Gamma = c_j^{\text{inv}} \cdot \sum_j \sum_v x_{j,v,T}^{\text{vintage}} \cdot \max\left(0, \frac{\left\lceil \frac{L_j}{2} \right\rceil - \left(\frac{T}{2} - \left\lceil \frac{v}{2} \right\rceil\right)}{\left\lceil \frac{L_j}{2} \right\rceil}\right), \tag{5}$$

where T is the total number of stages. The numerator of the last term of (5) is the remaining lifetime of technology j beyond the planning horizon; hence, the last term is the fraction of the remaining lifetime. The investment subproblem, SP_{inv}, is formulated as:

$$\min\ \alpha^{\lceil t/2 \rceil - 1} \sum_j \left(c_j^{\text{inv}} \cdot u_{j,t}^{\text{inv}} + c_j^{\text{fix}} \cdot x_{j,t}^{\text{alive}}\right) - \Gamma, \tag{6}$$

subject to:

$$x_{j,t,t+1}^{\text{vintage}} = x_{j,t,t}^{\text{vintage}} + u_{j,t}^{\text{inv}}, \qquad t \in \mathcal{T}_{\text{inv}}, \forall j, \tag{7}$$
$$x_{j,v,t+1}^{\text{vintage}} = x_{j,v,t}^{\text{vintage}}, \qquad \forall v \neq t, \forall j, \tag{8}$$
$$x_{t+1}^d = x_t^d, \qquad t \in \mathcal{T}_{\text{inv}}, \tag{9}$$
$$x_{j,0,1}^{\text{vintage}} = c_j^{\text{init}}, \qquad \forall j, \tag{10}$$
$$x_{j,v,1}^{\text{vintage}} = 0, \qquad \forall v \neq 0, \tag{11}$$
$$x_0^d = 1, \tag{12}$$
$$u_{j,t}^{\text{inv}} \geq 0, \qquad t \in \mathcal{T}_{\text{inv}}, \forall j, \tag{13}$$

where $u_{j,t}^{\text{inv}}$ is the new investment in technology j at time t, c_j^{inv} refers to investment cost per unit capacity, c_j^{fix} represents fixed annual operating costs, α is the discount factor, and c_j^{init} is the existing capacity of technology j at the beginning of the planning horizon.

C. Operational Subproblem

At each operational stage $t \in \mathcal{T}_{\text{opr}}$, we solve the following subproblem to determine the optimal hourly dispatch. Let $u_{j,t,h}^{\text{prod}} \geq 0$ be the thermal generation of technology j in hour h, and let $u_{t,h}^{\text{unmet}} \geq 0$ represent any unmet demand. The energy price ω_t and the annual demand multiplier ξ_t are realized at the start of this stage. The problem, denoted by SP_{opr}, is formulated as:

$$\min\ \mathbb{E}_{\omega_t}\left[\alpha^{\lceil t/2 \rceil - 1} \cdot \sum_j \sum_h \left(\left(\omega_t + c_t^{CO_2} \cdot c_j^{\text{emis}}\right) \cdot \frac{u_{j,t,h}^{\text{prod}}}{\eta_j} + c_j^{\text{var}} \cdot u_{j,t,h}^{\text{prod}} + c^{\text{penalty}} \cdot u_{t,h}^{\text{unmet}}\right)\right], \tag{14}$$

subject to :

$$\sum_j u_{j,t,h}^{\text{prod}} + u_{t,h}^{\text{unmet}} = c_h^d \cdot x_t^d \cdot \xi_t, \qquad \forall t \in \mathcal{T}_{\text{opr}}, \forall h, \tag{15}$$
$$x_{t+1}^d = x_t^d \cdot \xi_t, \qquad \forall t \in \mathcal{T}_{\text{opr}}, \tag{16}$$
$$u_{j,t,h}^{\text{prod}} \leq x_{j,t}^{\text{alive}}, \qquad \forall t \in \mathcal{T}_{\text{opr}}, \forall j, \forall h, \tag{17}$$
$$x_{j,v,t+1}^{\text{vintage}} = x_{j,v,t}^{\text{vintage}}, \qquad \forall t \in \mathcal{T}_{\text{opr}}, \forall j, \forall v, \tag{18}$$
$$u_{j,t,h}^{\text{prod}}, u_{t,h}^{\text{unmet}} \geq 0, \qquad \forall t \in \mathcal{T}_{\text{opr}}, \forall j, \forall h. \tag{19}$$

In (14), α is the discount factor, ω_t is the realized energy price, c^{CO_2} is the carbon price, and c_j^{emis} is the CO_2 emission factor of technology j. The term $\frac{u_{j,t,h}^{\text{prod}}}{\eta_j}$ reflects the required energy use to produce $u_{j,t,h}^{\text{prod}}$ MWh of heat, given thermal efficiency or the seasonal coefficient of performance in the case of heat pumpsη_j. A high penalty c^{penalty} discourages unmet demand. c_h^d is the base hourly load, scaled by the demand multiplier x_t^d and the annual demand multiplier ξ_t to yield the demand in hour h. Finally, the state variable for demand updates to $x_{t+1}^d = x_t^d \cdot \xi_t$. This ensures that future stages see the updated annual demand multiplier, consistent with the Markovian realisation ξ_t.

III. Case Study Setup

We implement the model for a hypothetical Central European city over four model years: 2020, 2030, 2040, and 2050, each representing the following ten-year period. In this setup, investment and operational decisions are made in the first year of each ten-year period. The operational costs are then scaled up by ten to represent the ten-year period. All the costs are discounted to the beginning of the first year, 2020.

We consider four candidate DH supply technologies: natural gas boiler and Combined Heat & Power (CHP), large-scale heat pump, and a deep geothermal (direct use) plant. Table I lists each technology's initial capacity, maximum additional capacity, thermal efficiency, and lifetime. The key cost and emission data for the technologies is available in Table II and in Table III in Appendix A. Please note that, in the current model setup, we only implement a carbon price to phase out fossil fuels and do not force an emission limit. The carbon price assumptions are given in Table V in Appendix A.

TABLE I
Technology parameters.

Technology	Initial Capacity	Max Expansion	Thermal Efficiency	Lifetime
Gas CHP	500 MW_{th}	500 MW_{th}	0.44	30 years
Gas Boiler	350 MW_{th}	500 MW_{th}	0.90	30 years
Heat Pump	-	250 MW_{th}	3.00	20 years
Geothermal	-	100 MW_{th}	1.00	30 years

Annual demand evolves through three states, $\xi \in \{1.1, 1.0, 0.9\}$, representing a 10% increase, stable, or 10% decrease, respectively. These three states were selected to capture plausible 5-year variations in a hypothetical setting while limiting complexity. The base annual demand is 2 TWh with a peak load of 814 MW. We assume the following transition matrix:

$$\Phi_{i,j} = \begin{bmatrix} 0.2 & 0.5 & 0.3 \\ 0.2 & 0.5 & 0.3 \\ 0.2 & 0.5 & 0.3 \end{bmatrix},$$

governs upward, stable, or downward movement, while the first model year is deterministic in terms of annual demand. Therefore, the annual demand increases, stays the same, or decreases by the probabilities of $0.2, 0.5, 0, 3$, respectively. These probabilities reflect the likelihood that demand remains within a similar range while shifting upwards or downwards with lower probabilities.

We model the natural gas price using a log-normal distribution to capture right-skewed price fluctuations with a mean of 37€/MWh and a standard deviation of 10€/MWh, discretising by four quantile points. Each operational stage draws from these distributions independently. The hourly electricity prices are assumed to be deterministic. We use the electricity price data of [13], distinguished by two profiles: until and after 2030.

We approximate the load profile within each operational stage by selecting 100 typical hours, following the method proposed by [14], aggregated from a full 8760-hour load profile of [13]. This ensures computational tractability in the operational subproblem. Each typical hour is assigned a weight for how many actual hours it represents. The details of the typical hour approach are given in Appendix B.

IV. Results and Discussion

A total of 400 scenarios were generated to solve the model under different possible future states. Each scenario combines (1) a randomly sampled path of the Markov demand state over the planning horizon and (2) a corresponding draw from the log-normal distribution for each operational year. These independent draws yield a sufficiently large, yet computationally manageable, set of plausible future outcomes, enabling robust capacity expansion decisions.

A. Alive Capacities

Figure 2 and 3 illustrate the alive capacities over eight stages (four model years) for the boiler and the heat Pump, respectively. The median (orange line) and the 25–75 % and 5–95 % quantiles (shaded regions) show how capacity evolves across 400 simulated scenarios.

As seen in Figure 2, the model starts with a significant existing capacity (350 MW) of gas boilers but eventually retires a large portion around stages 3–4. In many scenarios, the capacity later rebounds to above 400 MW, reflecting reinvestments in new gas boiler units if they remain competitive versus rising demand and uncertain fuel prices. The wide band between the 25–75 % and 5–95 % quantiles indicates that some scenarios see minimal new gas Boiler adoption. In contrast, others return to high capacities, depending on how demand and natural gas prices evolve.

Figure 3 shows the minimal capacity of heat pumps in the early stages but a sudden jump near the mid-horizon. The results are driven by the lifetime which was assumed to be lower for heat pumps than for gas boilers and geothermal capacities. The model's branching scenarios produce high uncertainty bands in the final two stages, indicating that some outcomes favour continued heat pump use while others retire it. Uptake of geothermal in the last phase also drives the lower role of heat pumps.

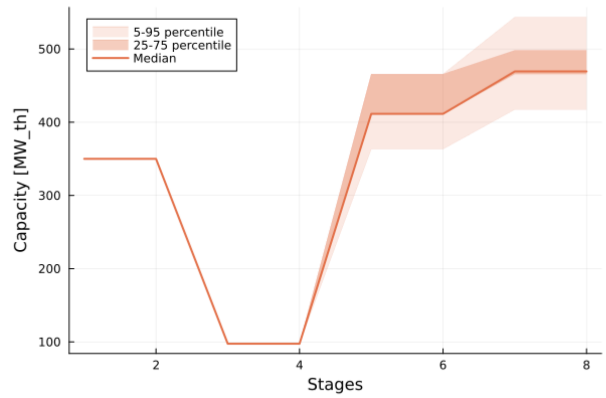

Fig. 2. Gas boiler alive capacity across eight all stages.

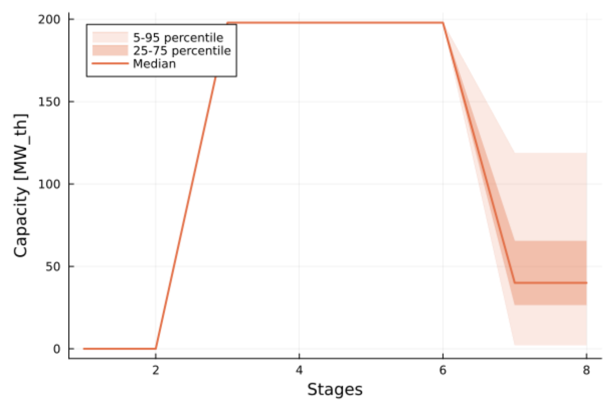

Fig. 3. Heat pump alive capacity across all stages.

The alive capacity plots for CHP and geothermal are given in Figure 7 and Figure 8 in Appendix C. CHP alive capacity begins around 500 MW but eventually drops near 100 MW or lower by the final stages in most scenarios, suggesting that high future carbon costs penalise gas CHP. Meanwhile, geothermal alive capacity grows from 0 to 300 MW by the last year in all scenarios, reflecting low marginal costs.

B. Investment Decisions

Figure 4 and 5 focus on newly built capacity (MW_th) at each investment stage $t \in \{1, 3, 5, 7\}$. We again plot the median and quantile bands across the stochastic scenarios.

The largest median investment in gas boilers often occurs around stage 3 (i.e. 2030 if each stage represents 10 years), where scenarios place 200–300 MW of new boiler capacity. Investment drops markedly by stage 5, and by stage 4 no investments happen due to the investment lead time. The high

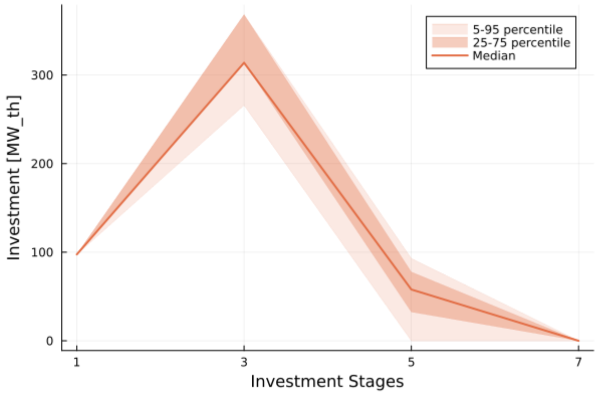

Fig. 4. Boiler investment decisions at four investment stages.

spike around stage 3 is consistent with retiring older boilers and installing new ones early, before gas and carbon prices grow too prohibitive.

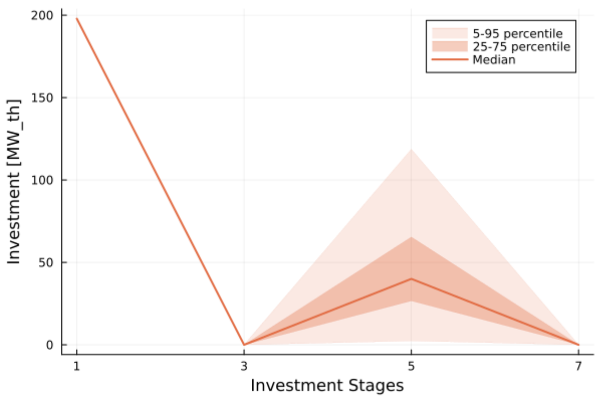

Fig. 5. Heat Pump investment decisions at four investment stages.

At stage 1, some scenarios invest as much as 200 MW in heat pumps, anticipating favourable future conditions. The median at stage 3 then drops sharply to zero, but re-emerges around stage 5 for a subset of scenarios. Finally, by stage 7, no expansions occur because of the investment lead time.

The investment into heat pump and CHP is given in Figure 10 and Figure 9 in Appendix C. CHP Investment sees a moderate spike at stage 3 but quickly vanishes to zero in later stages, consistent with carbon price escalation. Geothermal shows a consistent investment of 100 MW in all stages, given its low marginal cost and risk-free operational cost structure since it is not dependent on a fuel.

C. Operation of Technologies

Figure 6 in Appendix C presents violin plots of each technology's annual production. The vertical spread of each violin shape illustrates the range of production across different stochastic realizations (demand states, price draws). For example, in the boiler plot, we see a cluster of production levels near 2.0×10^6 MWh in the early years, but that distribution narrows significantly by 2050, indicating fewer scenarios where the Boiler remains the cheapest dispatch option. Meanwhile, Heat Pumps show almost the opposite pattern: very low production in 2030 but a wide distribution in 2050 and 2060, reflecting the model's uncertainty about how beneficial heat pumps can become under different price and demand realizations.

These violin plots thus highlight not only the expected production levels but also the risk or variability in each technology's utilization, which is an essential piece of information for decision-makers who seek robust, long-term planning strategies. Supplementary explanation on how the dispatch of the technologies look like is provided in Appendix D.

D. Discussions

Overall, these results confirm that a multistage approach captures significant re-timing of investments under uncertainty. For example, the model invests in gas boilers early when gas is relatively cheap and electricity prices are volatile, while concurrently adding some heat pumps or CHP if carbon pricing or operational synergy is attractive. Over time, zero-emission or electrified options (like geothermal) expand under strict carbon prices. Although these findings hold under our current cost structures and assumed geothermal availability, outcomes could shift if availability of the heat sources or technology costs differ. Furthermore, the salvage value term at the horizon end prevents the model from underestimating the value of new long-lived units toward the final stage. Still, short-lifetime or flexible technologies remain favourable in scenarios with high uncertainty, where option value is critical.

V. Conclusions

This paper proposed and tested a multistage stochastic optimization framework for DH capacity expansion, highlighting how separate investment and operational stages, Markovian demand evolution, and stagewise-independent fuel prices can be integrated into an SDDP model. The numerical illustrations show that boiler capacity can drop and rebound in response to mid-horizon retirements. Heat pumps, which are assumed to have a shorter lifetime than the other technologies, appear in some scenarios once energy prices are favourable, then are retired if those conditions deteriorate. CHP capacity is largely phased out as carbon costs escalate, with minimal reinvestment. Geothermal capacity expands steadily, offering zero marginal emissions and stable operation once installed.

These insights underscore the importance of modelling investment timing over multiple stages under uncertainty. Future work may integrate thermal storage, decommissioning decisions, and risk measures to further explore cost-optimal decarbonization pathways for DH.

References

[1] European Commission. Recast energy efficiency directive proposal. Available online: https://energy.ec.europa.eu/topics/energy-efficiency/targets-directive-and-rules/recast-energy-efficiency-directive, 2022.

[2] Ilias Dimoulkas and Mikael Amelin. Monte Carlo simulation of district heating system short-term operation in electricity markets. *Energetika*, 63(3), November 2017. Number: 3.

[3] Thomas T. D. Tran and Amanda D. Smith. Stochastic Optimization for Integration of Renewable Energy Technologies in District Energy Systems for Cost-Effective Use. *Energies*, 12(3):533, January 2019. Number: 3 Publisher: Multidisciplinary Digital Publishing Institute.

[4] Puneet Saini, Pei Huang, Frank fiedler, Anna Volkova, and Xingxing Zhang. Techno-economic analysis of a 5th generation district heating system using thermo-hydraulic model: A multi-objective analysis for a case study in heating dominated climate. *Energy and Buildings*, 296:113347, October 2023.

[5] Annelies Vandermeulen, Ina De Jaeger, Tijs Van Oevelen, Dirk Saelens, and Lieve Helsen. Analysis of Building Parameter Uncertainty in District Heating for Optimal Control of Network Flexibility. *Energies*, 13(23):6220, January 2020. Number: 23 Publisher: Multidisciplinary Digital Publishing Institute.

[6] Andreas Bott, Tim Janke, and Florian Steinke. Deep learning-enabled MCMC for probabilistic state estimation in district heating grids. *Applied Energy*, 336:120837, April 2023.

[7] Marc Hohmann, Joseph Warrington, and John Lygeros. A two-stage polynomial approach to stochastic optimization of district heating networks. *Sustainable Energy, Grids and Networks*, 17:100177, March 2019.

[8] Amos Schledorn, Daniela Guericke, Anders N. Andersen, and Henrik Madsen. Optimising block bids of district heating operators to the day-ahead electricity market using stochastic programming. *Smart Energy*, 1:100004, 2 2021.

[9] Xin Shen, Zhao Luo, Jun Xiong, Hongzhi Liu, Xin Lv, Taiyang Tan, Jianwei Zhang, Yuting Wang, and Yinghao Dai. Optimal Hybrid Energy Storage System Planning of Community Multi-Energy System Based on Two-Stage Stochastic Programming. *IEEE Access*, 9:61035–61047, 2021. Conference Name: IEEE Access.

[10] Oscar Dowson. The policy graph decomposition of multistage stochastic programming problems. *Networks*, 76(1):3–23, 2020.

[11] J. Hole, A. B. Philpott, and O. Dowson. Capacity planning of renewable energy systems using stochastic dual dynamic programming. *European Journal of Operational Research*, 322(2):573–588, April 2025.

[12] O. Dowson and L. Kapelevich. SDDP.jl: a Julia package for stochastic dual dynamic programming. *INFORMS Journal on Computing*, 33:27–33, 2021.

[13] Mostafa Fallahnejad, Marcus Hummel, Daniel Keshaw, Andreas Müller, Ardak Akhatova, Ali Kök, Lukas Kranzl, and Aadit Malla. Umfassende bewertung des potenzials für eine effiziente wärme- und kälteversorgung [Comprehensive Assessment of the Potential for Efficient Heating and Cooling]. Technical report, Bundesministerium für Klimaschutz, Umwelt, Energie, Mobilität, Innovation und Technologie, Vienna, Austria, 2024.

[14] P. Gabrielli, M. Gazzani, E. Martelli, and M. Mazzotti. Optimal design of multi-energy systems with seasonal storage. *Applied Energy*, 219:408–424, June 2018.

APPENDIX

A. Cost and Emission Parameters

Table II shows the principal cost parameters of the candidate technologies.

TABLE II
COST PARAMETERS

Technology	Investment Cost	Fixed O&M	Variable O&M
CHP	1.1 M€/MW	15 k€/MW/a	3.3 €/MWh
Boiler	0.1 M€/MW	2 k€/MW/a	0.75 €/MWh
HeatPump	0.5 M€/MW	3 k€/MW/a	1.1 €/MWh
Geothermal	1 M€/MW	2 k€/MW/a	0.5 €/MWh

Table III shows the energy carriers and the respective emission factors of the candidate technologies.

TABLE III
EMISSION AND ENERGY CARRIER PARAMETERS

Technology	Emission Factor	Energy carrier
CHP	0.2 tCO_2/MWh	Natural gas
Boiler	0.2 tCO_2/MWh	Natural gas
Heat Pump	-	Electricity
Geothermal	-	Geothermal heat

Table IV outlines the discretisation approach for stochastic natural gas prices, following a log-normal distribution. The probabilities in Table IV were derived by discretizing this continuous distribution at quantiles (e.g., 0.05, 0.27, 0.5, 0.73, 0.95), ensuring each discrete price node represents a particular percentile of the underlying distribution.

TABLE IV
LOG-NORMAL PRICE DISCRETIZATION

Quantile	Gas Price	Probability
5%	23.08	0.16
35%	32.25	0.42
65%	39.57	0.35
95%	55.28	0.07

Table V shows the evolution of the carbon price over the modelling horizon.

TABLE V
CO_2 PRICE EVOLUTION OVER THE YEARS

Period	CO_2 Price
2020-2029	50 €/tCO_2
2030-2039	150 €/tCO_2
2040-2049	250 €/tCO_2
2050-2059	350 €/tCO_2

B. Typical Hours

The number of hours over a year is 8760. In practice, optimisation models of this kind use typical weeks/days/hours to reduce the complexity of the problem while maintaining an acceptable accuracy. Ref [14] show how decreasing the number of typical days reduces the accuracy of their model but also accelerates the obtention of acceptable results.

In this case, as there is no interdependence between timesteps (absence of storage, inertia, etc.), typical hours are considered to represent our problem. To compute them, we perform a stepwise approximation considering N typical loads or typical hours. To do so, we consider N equally distant quantiles of the original load, starting from 0% to 100%. Then, we assign each original load to the closest typical load as an approximation. In this article, the considered number of typical hours is $N = 100$.

C. Additional Results Figures

We provide violin plots of each technology's annual production, alive capacity and investment figures for CHP and geothermal that were omitted from the main text.

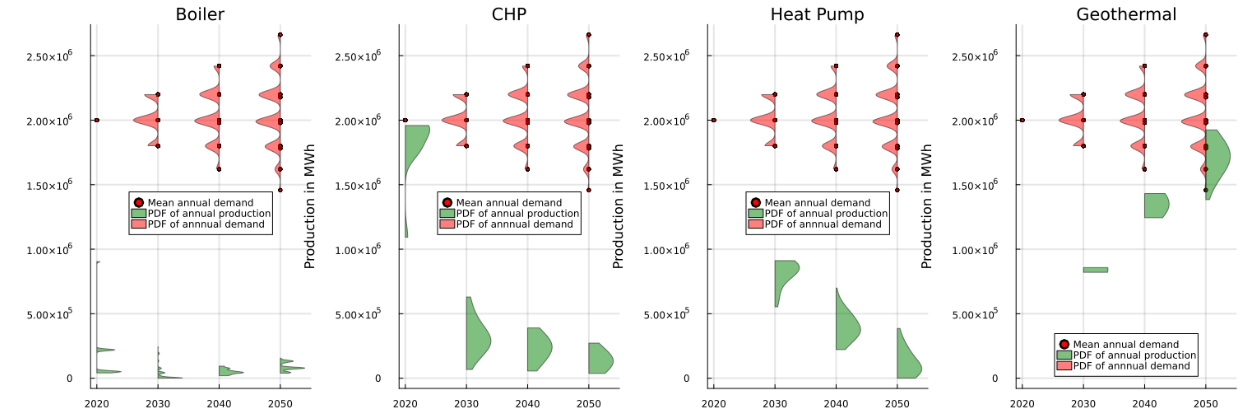

Fig. 6. Violin plots of each technology's annual production in MWh at the four operational stages (2020, 2030, 2040, 2050 on the x-axis), overlaid with the annual demand of the entire DH system

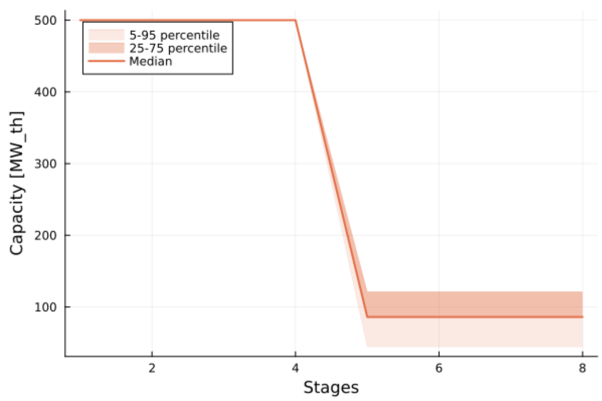

Fig. 7. CHP alive capacities across all stages.

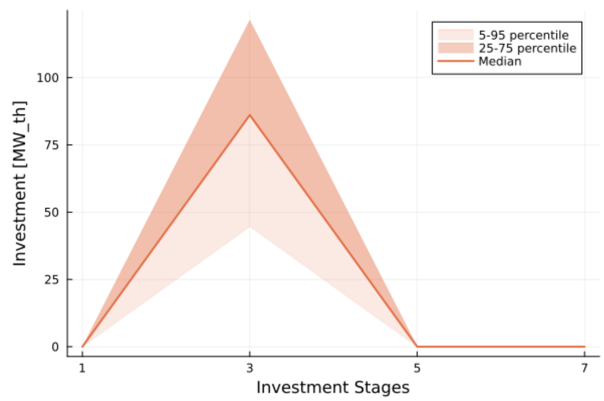

Fig. 9. CHP investment decisions at investment stages.

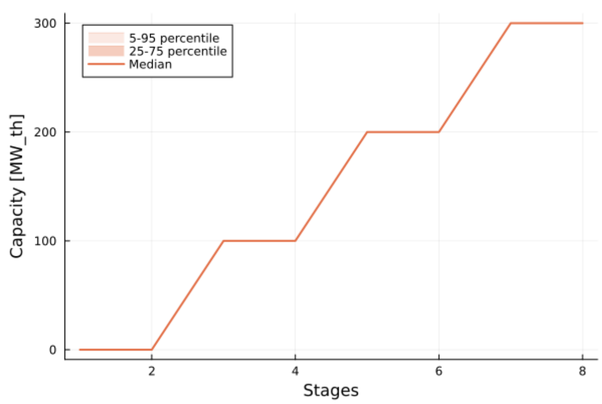

Fig. 8. Geothermal alive capacities across all stages.

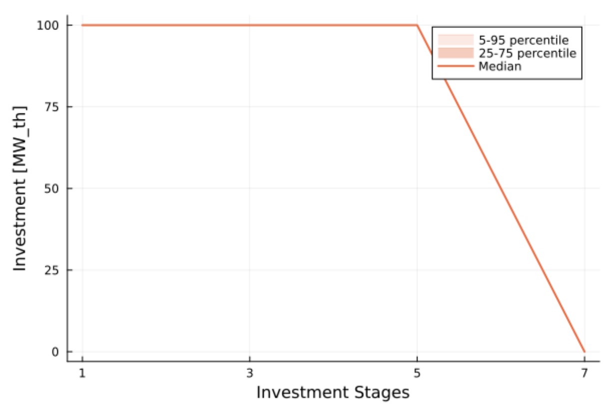

Fig. 10. Geothermal investment decisions at investment stages.

D. Dispatch of technologies

To illustrate how these capacities operate during each model year, Figure 11 shows a sample load duration curve decomposition for one representative scenario. The horizontal axis accumulates the hours of each operational stage (e.g. 8760 hours $\times$ four years), sorted by demand. The stacked area indicates how the dispatch is split among CHP, Heat Pump, Boiler, and Geothermal. We see that in earlier stages, the model relies heavily on CHP for base load, with the boiler covering residual peaks. Heat pumps and geothermal appear in the mid-horizon stages, with heat pumps covering moderate segments of the load curve, effectively displacing the Boiler. Geothermal covers a relatively small fraction of the baseload, reflecting its lower marginal cost but higher investment cost

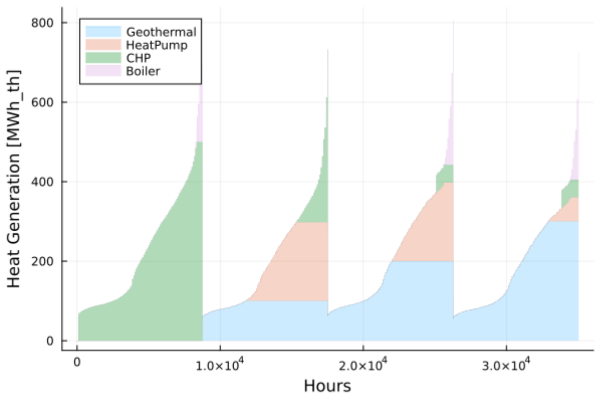

Fig. 11. Load duration curve of a sample simulation over the hours of operational stages.

and constrained capacity potential. In the final model year (rightmost portion of the figure), CHP capacity ramps up for peak segments while Heat Pump usage declines in some scenarios, consistent with the notion of short-lifetime heat pump units being retired and no longer needed due to the increased capacities of geothermal.