Sequential Investment Optimization for District Heating Capacity Expansion

Ali Kök[(1)], Jonathan Hachez[(2)], Lukas Kranzl[(1)], Shravan Kumar[(3)], Johan Dalgren[(4)]
[(1)]Energy Economics Group, TU Wien
[(2)] Vrije Universiteit Brussel
[(3)] KTH Royal Institute of Technology
[(4)] Stockholm Exergi AB

EPOC Winter Workshop, 4 September 2026

Outline

- **Motivation and Background**
  - What district heating is
  - The planning challenge
- **Literature, gap and research questions**
- **Methodology**
  - The decision problem in SDDP terms
  - Policy graph and formulation
  - Why SDDP
- **Case study: Stockholm**
  - Uncertainty as a Markov chain
  - Uncertainty parameterisation
- **Results**
  - Investment decisions and capacity evolution
  - Value of the stochastic solution
  - Convergence
- **Conclusions and the EPOC project**

What district heating is

Central heat production, a hot water network, and the buildings it serves

- Heat is produced in central plants and delivered as hot water through a pipe network to buildings. One operator plans production, network and supply contracts together.
- Fourth generation DH (4GDH) lowers supply temperatures and runs on heat pumps, which couples the system to the electricity market and to building retrofits.

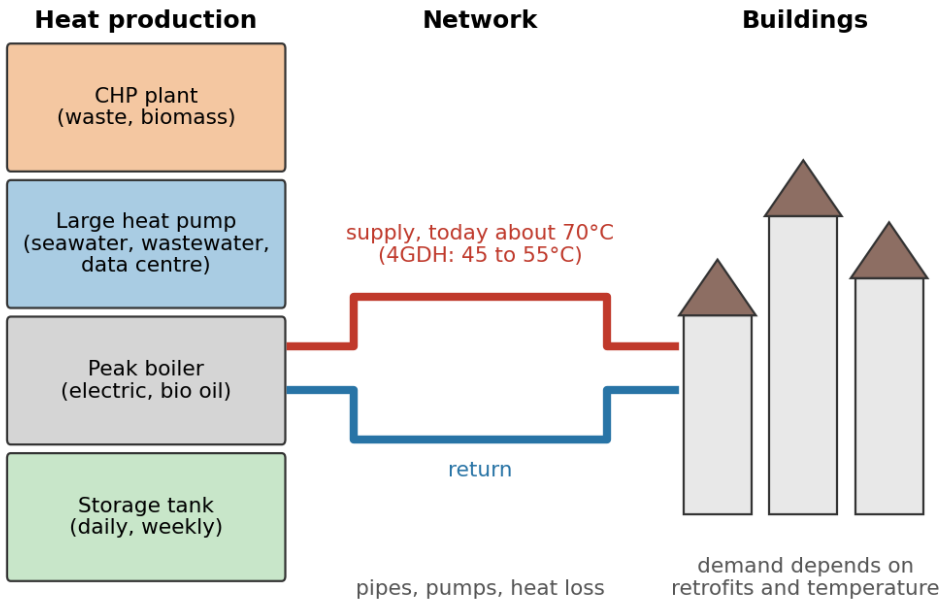

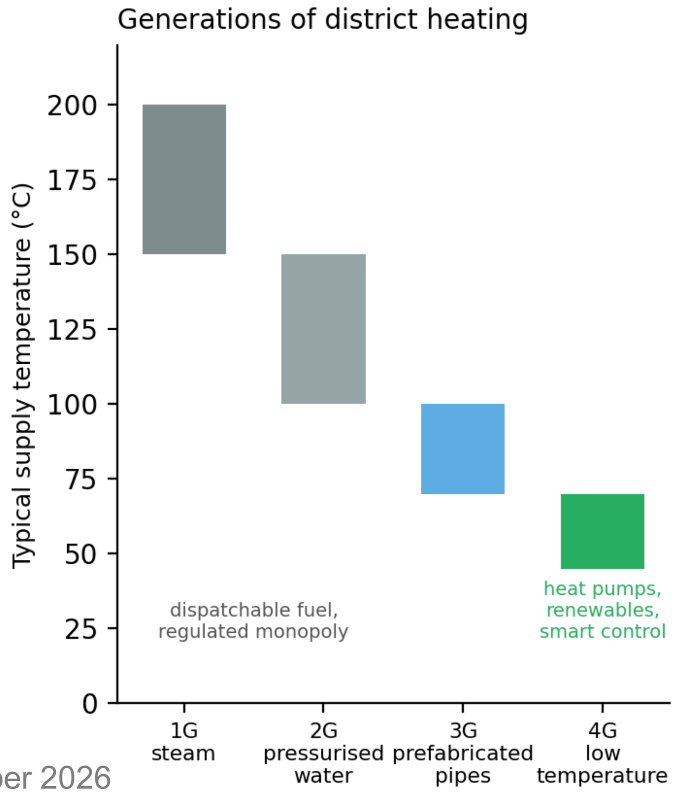

The Planning Challenge

DH operators face irreversible investment decisions under deep uncertainty → risking stranded assets

- Investment decisions involve hundreds of millions EUR → cannot be reversed once commissioned
- Uncertainties have grown significantly:
  - 2022 energy crisis: prices increased dramatically within months (IEA, 2022)
  - EU ETS carbon prices: <€10 (2018) → >€100 (2023) → €150–200 by 2030 (Pietzcker et al., 2021)
- Structural changes: heat demand evolution, transition to low-temperature DH
- These uncertainties interact → fuel and power prices move together, and the network's supply temperature changes both heat pump efficiency and demand

District heating, for an optimisation audience

Central heat production, no reservoir, one slow structural transition

- City-scale hot water network. Stockholm: 800,000 people, about 8 TWh heat per year
- CHP plants burn fuel and sell electricity; heat pumps buy electricity (COP = heat out per electricity in); boilers cover peaks
- Storage: daily and weekly water tanks only
- Structural uncertainty: supply temperature. 70°C to 45°C raises COP 30 to 55 percent, needs retrofits that cut demand 20 percent
- Brownfield: 3,570 MW retires on a fixed schedule, most of it before 2040

Literature Review: DH Capacity Planning

A dichotomy exists: stochastic operations vs. deterministic capacity planning

- **Short-term operational planning (hours to days):**
  - Stochastic programming is mature (Schledorn et al., 2021; Guericke et al., 2024); Multistage for HP dispatch (Siddique et al., 2024)
- **Long-term capacity expansion (decades):**
  - Deterministic MILP dominates (Kök et al., 2025; Kumar et al., 2025, …)
  - Two-stage stochastic at best (Zwickl-Bernhard et al., 2025, …)
  - Lambert et al., 2016: multistage SP for network pipes → generation capacity fixed
- **Broader energy systems**: SDDP standard in hydro(power) (Dowson, 2020; Hole et al., 2025)
- **Liang et al., 2025: 109 studies reviewed → multistage SP and SDDP absent from DH capacity expansion literature**

Gap, contributions and research questions

What DH capacity planning lacks, and the questions this paper answers

- Three contributions
  - SDDP for DH generation capacity expansion
  - hybrid uncertainty: Markov price regimes plus a structural temperature branch
  - vintage tracking for brownfield systems (capacity tracked by installation year)
- Research questions
  - RQ1: How much does perfect foresight bias technology choice?
  - RQ2: What is the value of the stochastic solution?
  - RQ3: How does structural uncertainty about the temperature pathway shape investment timing and technology choice?

The decision problem in SDDP terms

Alternating investment and operational nodes, 55 states, no stagewise noise

- State: capacity by technology and vintage (10 x 5) plus storage vintages (5) = 55
- Controls: continuous investment; hourly dispatch over 4 representative weeks x 168 h
- Randomness: 3-state price Markov chain and a binary temperature branch, all in the graph
- 12 stages: 6 model years of 5 years, each an investment node followed by an operational node, 2025 to 2050, one model year build lead time
- Unlike hydro: the coupling across stages is capacity and asset age, not stored water

Policy Graph Structure

Decision hierarchy: alternating investment & operational stages

- Stage 1 (here-and-now): Investment decision without observing uncertainty
- Stage 2: Energy price uncertainty revealed (3 Markov states: H, M, L)
- Stage 3: Temperature pathway revealed → branching into High-temp / Low-temp
- Stages 4+: Markov price transitions within each temperature pathway

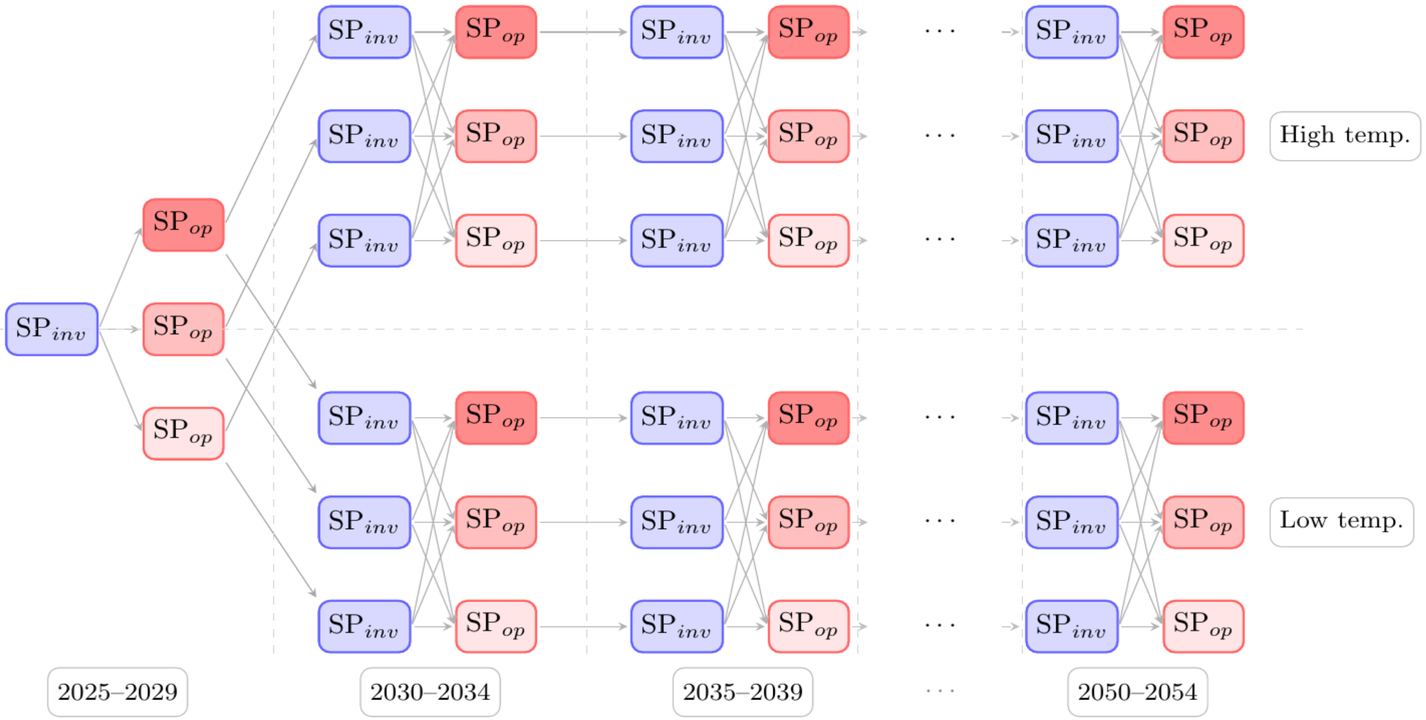

Results in 64 nodes, 55 states per node
1,458 paths over the 2025–2050 horizon

Formulation

Every subproblem is an LP

$$V_n(\mathbf{x}^{\text{in}}, \omega_n) = \min_{\mathbf{u}_n, \mathbf{x}^{\text{out}}} \left\{ C_n(\mathbf{x}^{\text{in}}, \mathbf{u}_n, \mathbf{x}^{\text{out}}, \omega_n) + \sum_{m \in \mathcal{S}(n)} P_{nm} V_m(\mathbf{x}^{\text{out}}, \omega_m) \right\} \quad (1)$$

subject to feasibility constraints $(\mathbf{x}^{\text{out}}, \mathbf{u}_n) \in \mathcal{X}_n(\mathbf{x}^{\text{in}}, \omega_n)$. Here, $\mathbf{x}^{\text{in}}$ is the incoming state vector (vintage capacities), $\mathbf{u}_n$ are local control variables, $\mathbf{x}^{\text{out}}$ is the outgoing state vector, and $\mathcal{S}(n)$ denotes the set of successor nodes reachable from n, and P_{nm} represents the conditional probability of transitioning from node n to node m. The random variable ω_n represents the uncertainty realization at node n, capturing temperature regime and energy market conditions

- **Investment stage objective:**
  - Minimize: investment costs + fixed O&M − salvage value + E[cost-to-go]
- **Operational stage objective:**
  - Minimize: fuel costs + carbon costs − electricity revenue + E[cost-to-go | ω_t]
- **Vintage tracking:**
  - x_alive = Σ_v x_vintage · $\lambda_{j,v,t}$ + $E_{j,k}$ (new + legacy capacity)

Why SDDP and not the deterministic equivalent

The graph is small; the operational subproblems are not

- Each operational node: 9,530 variables, 10,812 rows (672 hourly demand balances, 672 storage dynamics, 55 state transitions, 8,741 capacity and rate limits)
- Unrolled scenario tree: 2,908 nodes, about 21 million variables and 24 million rows
- Policy graph: 64 subproblems, 60 to 80 minutes on 24 cores
- More regimes, pathways or weeks grow the graph linearly, the tree exponentially

Case Study: Stockholm DH System

One of Europe's largest DH networks: 800,000 residents, ~8 TWh/year

- **System characteristics:**
  - >90% renewable supply (waste incineration, biomass CHP, heat pumps)
  - Base demand: 7,902 GWh/year; planning horizon: 2025–2050
- **Existing capacity (2025): ~3,570 MW total**
  - Waste CHP: 334 MW | Wood Chip CHP: 522 MW | Bio Oil Boiler: 1,714 MW
  - Seawater HP: 320 MW | Wastewater HP: 220 MW | Pellet + Bio Oil CHP: 460 MW
- **Planning challenge:**
  - Substantial capacity approaching end-of-life

Data source: Kumar et al., 2025

Uncertainty as a Markov chain

Three price regimes with 0.90 persistence, one absorbing temperature branch

- Regime in {High, Medium, Low} moves fuel and electricity prices together
- P rows (H, M, L): (0.90, 0.09, 0.01), (0.05, 0.90, 0.05), (0.01, 0.09, 0.90)
- Transitions every five years, between operational stages
- Temperature pathway 50/50, revealed once before the 2030 investment; sets COP and demand for good
- Carbon price and waste gate fees deterministic

Uncertainty Parameterization

Two pathways, 50/50, spanning the plausible extremes of network evolution

- Temperature pathways (50%/50% probability):
  - High-Temp: ~70°C maintained; demand +12% to 8,850 GWh (Kumar et al., 2025)
  - Low-Temp: 4GDH, 45°C by 2050; demand −20% to 6,322 GWh; COP improves: Data centre HP 4.0→6.2; Seawater HP 3.6→4.9

- Low-Temp demand reduction rationale:
  - Light efficiency measures reduce Stockholm demand ~18% (Pasichnyi et al., 2019)
  - Sufficient for low-temp compatibility (Casamassima et al., 2024)

Results: Heat Pump Investments

Data Center HP emerges as a 'no-regret' technology across all pathways

- Data Center HP: 672 MW in both pathways
  → robust hedging asset
- Seawater HP: 807–874 MW

  → largest investment, price-elastic under Low-Temp
- Wastewater HP: ~200 MW by 2030

  → early, robust commitment

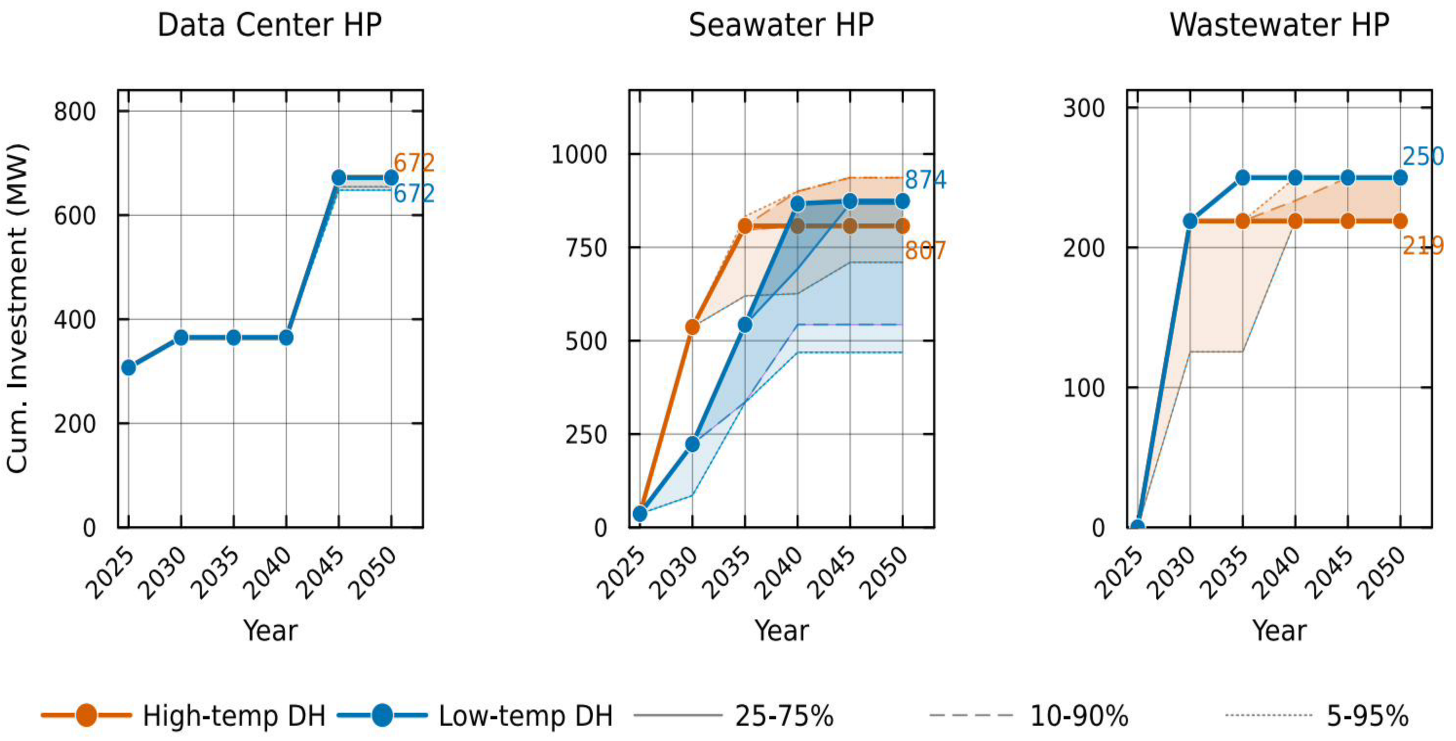

Results: CHP and Boiler Investments

Wood Chip CHP shows strongest pathway dependence: 804 MW vs. 0 MW

- **Wood Chip CHP:**
  - High-Temp: 804 MW (median) → baseload necessity
  - Low-Temp: 0 MW (median) → but retained as hedge in upper percentiles
- Waste CHP: 208–220 MW in both pathways
- Electric Boiler: Wide bands (up to 365 MW) → flexibility option
- Policy avoids biomass lock-in while retaining option value

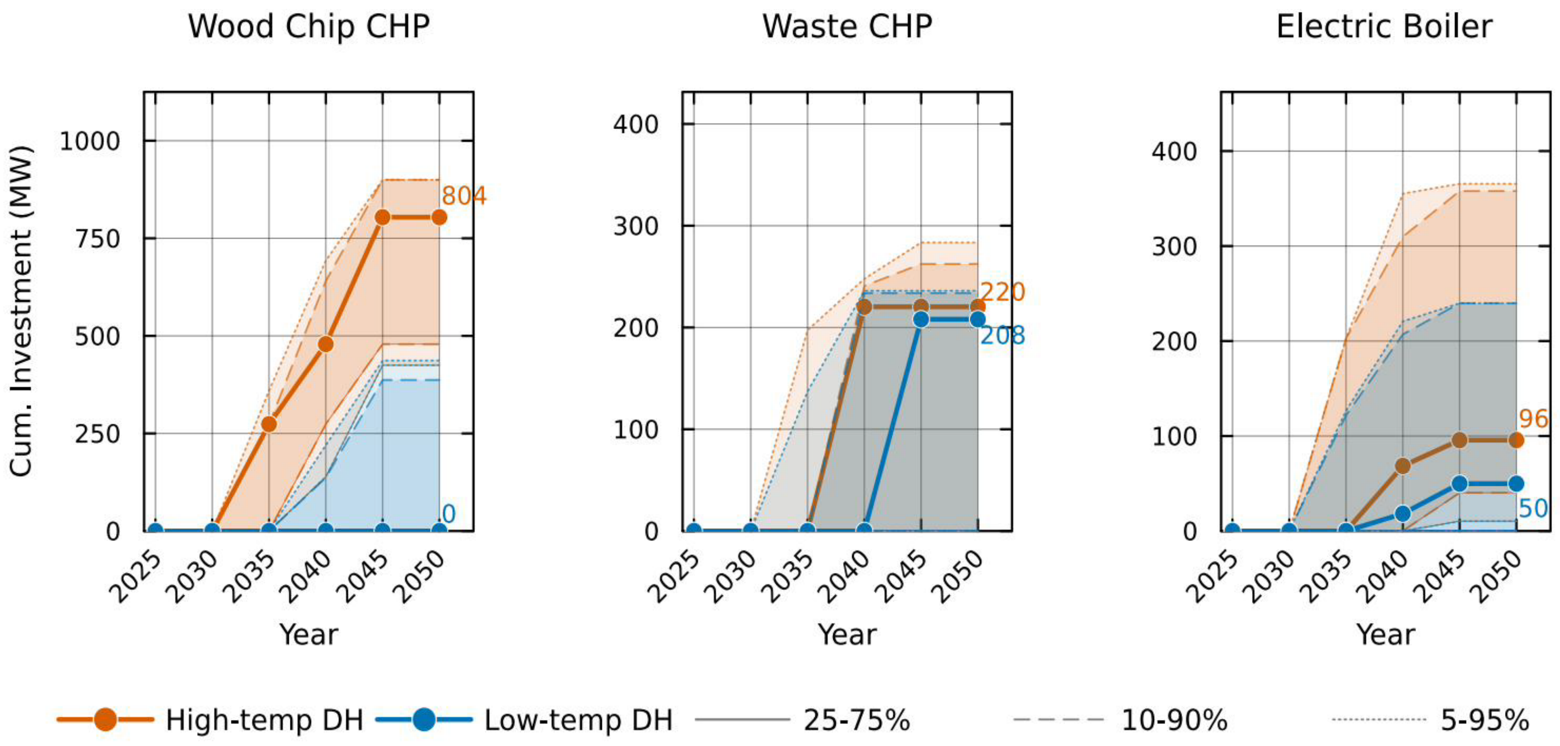

Results: Installed Capacity Evolution

Legacy bio-oil boilers bridge capacity gaps during system transition

- Heat pumps: Seawater HP shows a temporary dip (~2035) as old assets retire
- Brownfield flexibility: ~800 MW Bio Oil Boiler fleet bridges the gap
- Wood Chip CHP divergence:
  - High-Temp: 522 MW → 804 MW (growth)
  - Low-Temp: 522 MW → 0 MW (phase-out)
- Waste CHP: "V-shaped" trajectory → decline then selective reinvestment
- Policy delays capital-intensive replacements until strictly necessary

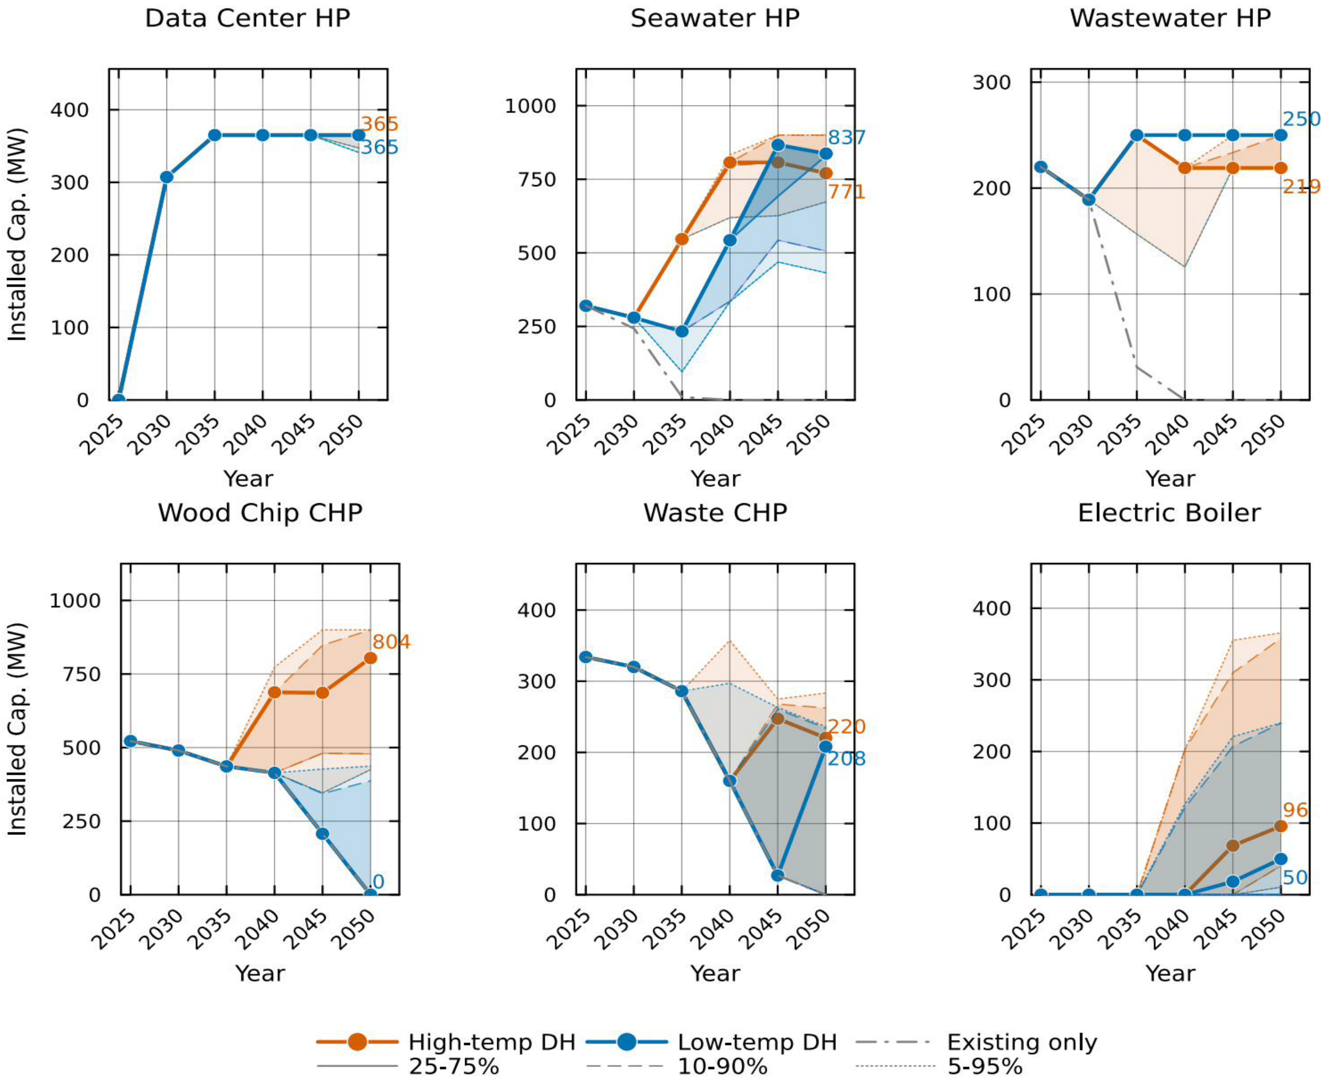

Results: Heat Generation Mix

Low-Temperature pathway enables complete inversion of the generation portfolio

- **High-Temperature pathway (2050):**
  - Total: 8,960 GWh (+13.4%); CHP: 62.5%; Heat pumps: 37%
- **Low-Temperature pathway (2050):**
  - Total: 6,404 GWh (−19.0%); Heat pumps: 73.8%; CHP: 26.1%
- Biomass CHP: 47.8% (High-Temp) vs. 6.0% (Low-Temp) → strategic divergence
- Residual biomass in Low-Temp mean (median is zero) = option value against price spikes

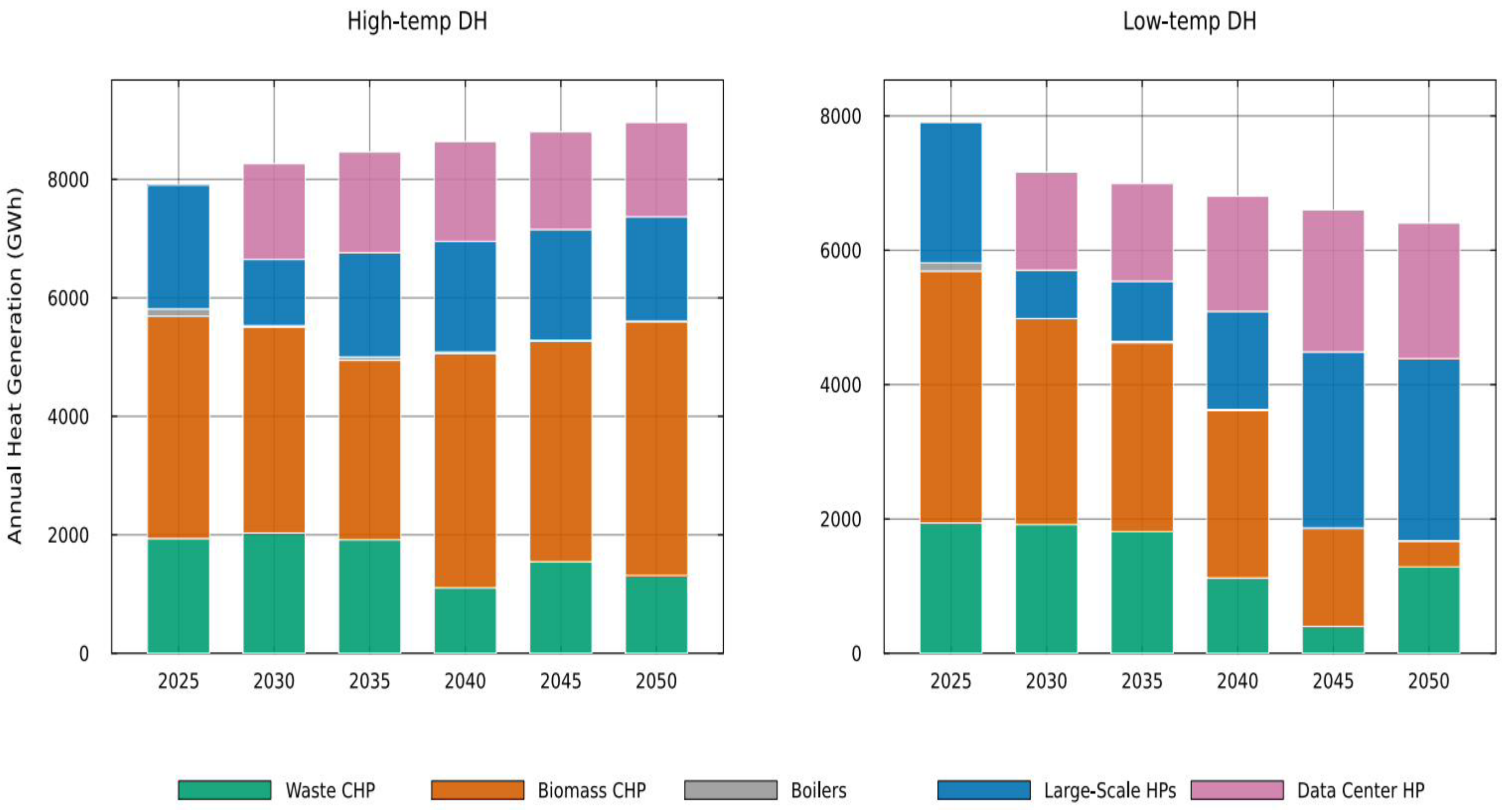

Value of the Stochastic Solution

VSS = 27% (risk-neutral) to 39% ($CVaR_{0.95}$) cost reduction

- SDDP policy: 7.39 billion EUR expected cost
- Deterministic (EV) policy: 10.15 billion EUR → VSS = 2.76 billion EUR (27.2%)
- VSS by pathway: High-Temp = 4.8 billion EUR (37%) vs. Low-Temp = 0.5 billion EUR (7%)
    - Deterministic under-invests in Stage 1 (2025: 226 MW vs 349 MW Data Center HP)
    - Stage 3 (2030): SDDP adds 500 MW Seawater HP vs. EV's 170 MW
  - When High-Temp materializes → 330 MW capacity gap → 4.8 billion EUR cost difference
- EEV: mean-value investment plan fixed at every stage, dispatch re-optimised, same 1,000 paths
- Curve: CVaR of each policy's simulated costs. The knee at α = 0.6 is where the worst outcomes become exactly the High-Temp branch

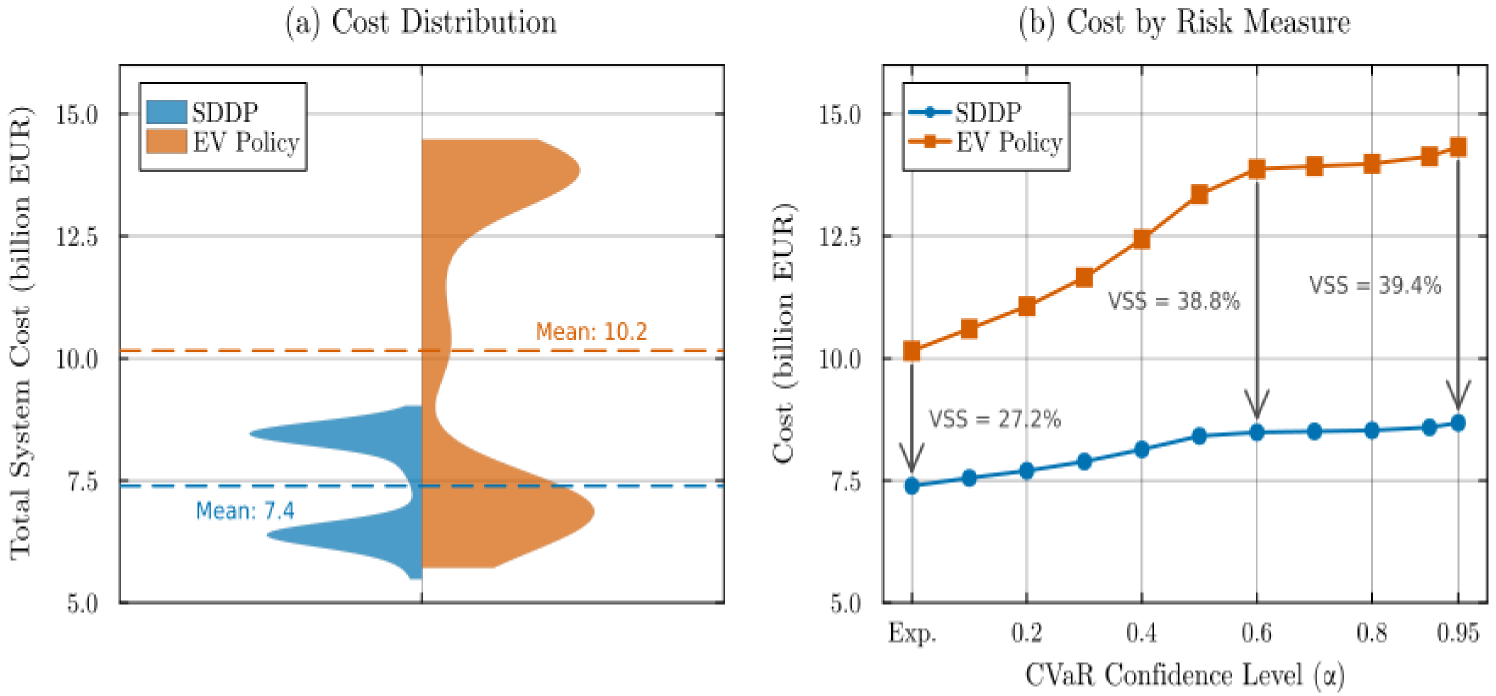

Convergence and computation

Final gap 0.5 percent, inside Monte Carlo error

- SDDP.jl v1.13, Gurobi 12, 24 threads
- Fixed iteration count, no gap criterion
- Lower bound 7.352 billion EUR, simulated cost 7.392 over 1,000 paths
- Bound flat to 0.14 percent after iteration 300
- Plot: rerun of the same model, 1,015 iterations

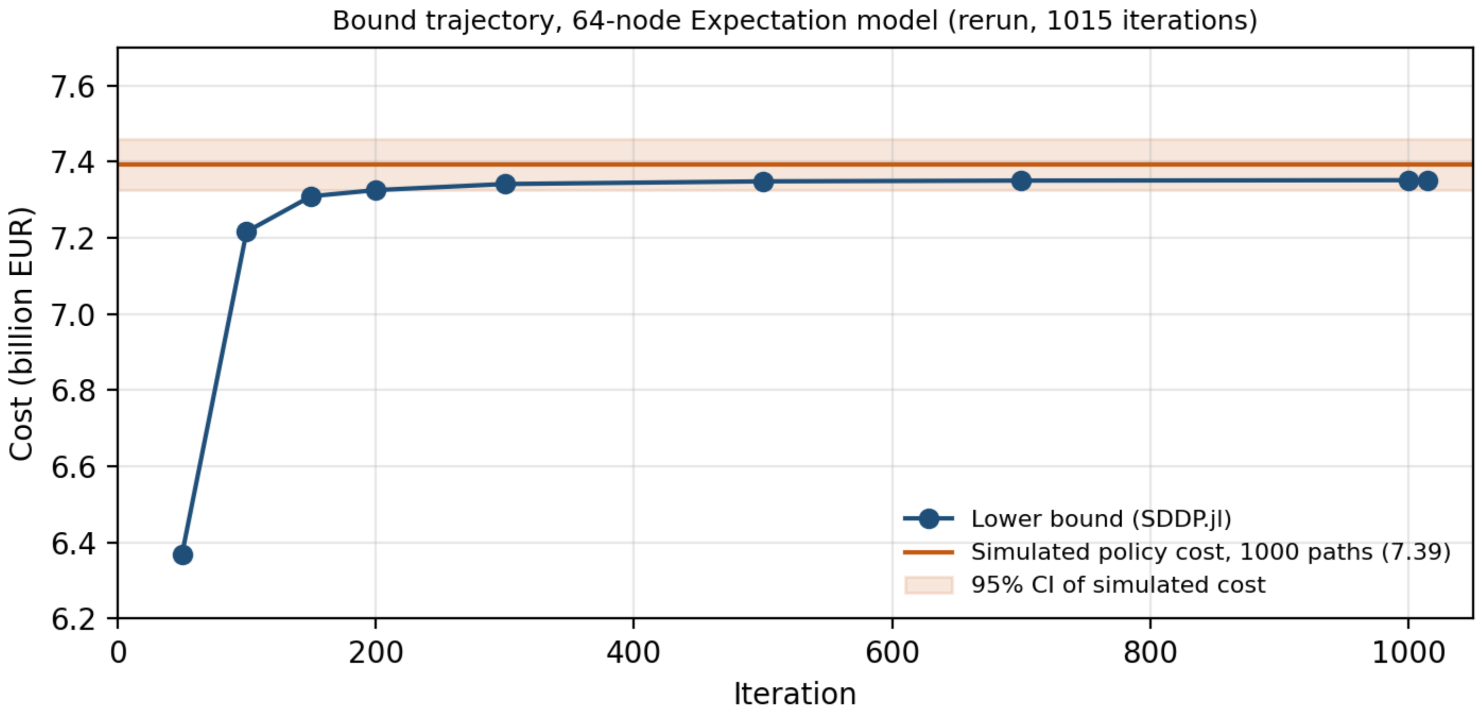

Conclusions

Multistage stochastic optimization is both tractable and valuable for DH

- RQ1 (Perfect foresight bias): Deterministic planning under-invests in foundational capacity → 330 MW gap → cost distribution 2x more volatile
- RQ2 (Value of Stochastic Solution): 27% risk-neutral, 39% under $CVaR_{0.95}$
- RQ3 (Decision hierarchy): No-regret technologies (Data Center HP, Wastewater HP) anchor first stage; biomass retains option value as hedge
- City-scale DH solved in about an hour with a 0.5% gap

What I am working on at EPOC

Coupling investment and operational planning for hydro-thermal systems

- The DH model nests investment nodes inside the SDDP graph. It works because the DH operational node is small
- Hydro operational models are much heavier; nesting them does not scale the same way
- Option A: a convex surrogate of operating cost in capacity, fitted from SDDP, embedded in JuDGE
- Option B: Epicycle graph → operations modelled as infinite horizon loops and shared between investment nodes
- Open questions: two-reservoir convex regression, convergence of the coupled iteration, validation against a full NZ model

Thank you for your attention! Questions?

koek@eeg.tuwien.ac.at

Kök, A., Hachez, J., Kranzl, L., Kumar, S., & Dalgren, J. (2026). Beyond perfect foresight: Sequential investment optimization for district heating capacity expansion. Energy Conversion and Management, 357, 121433. https://doi.org/10.1016/j.enconman.2026.121433

References

Backup: Legacy Capacity Retirement

Bio Oil Boiler fleet provides critical reserve capacity (~780 MW in 2035)

- Total legacy capacity: 3,570 MW (2025) → declining over horizon
- Bio Oil Boiler: 1,714 MW → ~780 MW (2035) → 0 MW (by ~2045)
- Wood Chip CHP: Complete phase-out by 2050
- **This creates the capacity gap that new investments must fill**

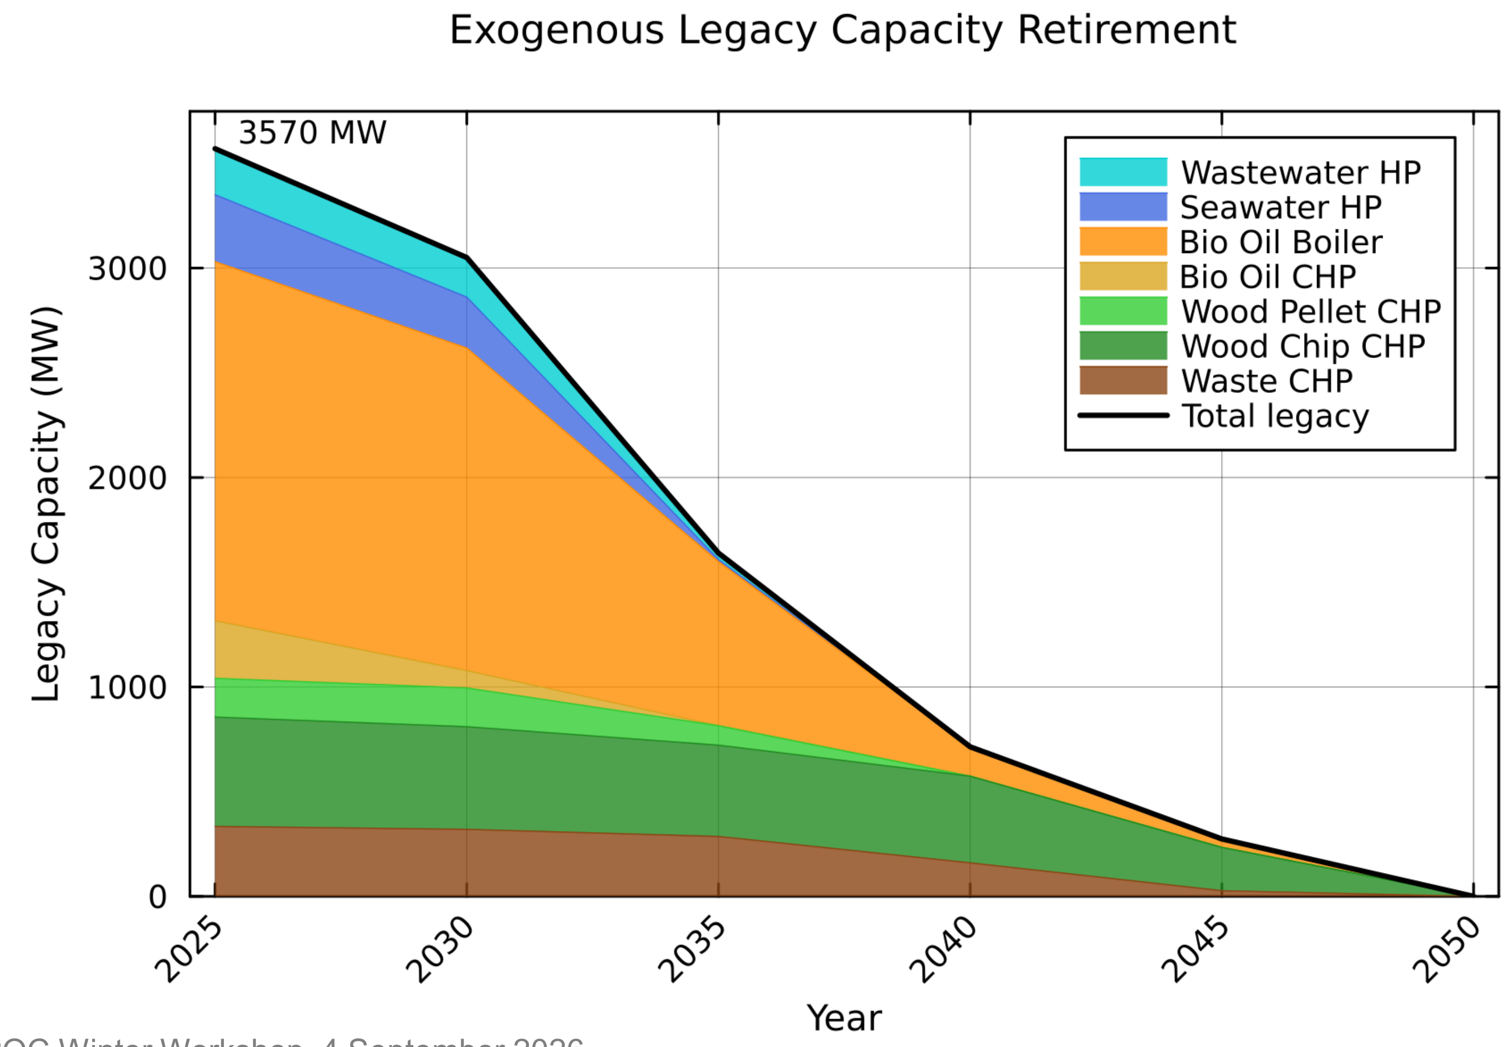

Backup: Supplementary Capacity Evolution

Technologies excluded from main analysis: Wood Pellet CHP, Bio Oil CHP/Boiler

- Wood Pellet CHP: 185 MW → 0 MW (retirement, no reinvestment)
- Bio Oil CHP: 275 MW → 0 MW (retirement, no reinvestment)
- Bio Oil Boiler: Critical reserve through 2035, then phase-out
- Low-Temp HP: Marginal investment in some scenarios