Navigating Energy Doldrums

Modeling the Impact of Energy Price Volatility on the HPC Cost of Ownership

Peter Arzt, Felix Wolf
TU Darmstadt

19th Workshop on Scheduling for Large Scale Systems
Fréjus, France

Motivation

- Fight against climate change
- Fossil energy is being phased out
- Renewables provide cheap, but intermittent alternatives

→Increased **energy price volatility**

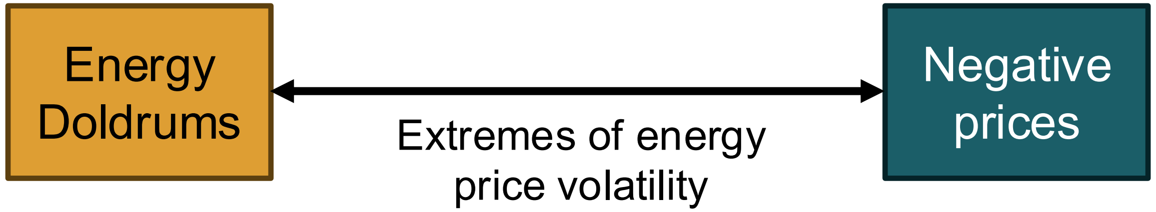

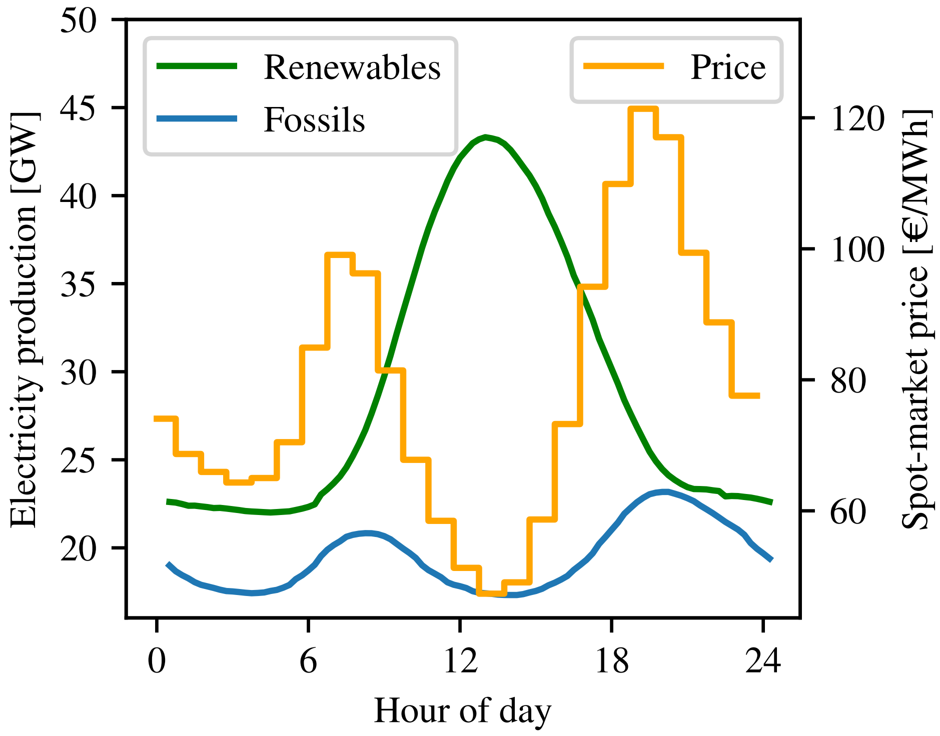

Electricity production and spot-market price in Germany for an average day in 2024

TECHNISCHE UNIVERSITÄT DARMSTADT
Parallel Programming

Variable Capacity

- Idea: Dynamically adjust capacity based on current prices
- Here: Temporary (partial) **shutdowns** of HPC cluster during high-price phases
- Introduces **trade-off** between energy cost savings and hardware underutilization

No shutdown

- 100% hardware utilization possible
- No energy savings

← Trade-off →

100% shutdown

- 0% hardware utilization
- No energy costs

Overview

Research questions:

- *Level of energy price fluctuation to make temporary shutdowns viable?*
- *Maximum cost-efficiency gains?*

This work:

Simple model to help assess trade-off

- Based on key system parameters
- Provides an upper bound for economic viability of temporary shutdowns
- Demonstrated in historical and hypothetical scenarios

Discussion of limitations and constraints

Categorizing Costs

For some time period of interest (e.g., 5 years):

- Total cost of ownership **TCO**
- Fixed costs **F**
  - Static costs that cannot be lowered by shutdowns
  - E.g., hardware write-offs, infrastructure, staff
- Energy costs **E**
  - Dynamic costs that can be reduced by shutdowns
  - I.e., electricity (mostly)

$$\mathbf{TCO = F + E}$$

Price Model

- *Electricity price forecasting* (EPF) is an active research field in economics (c.f., *Lago et al.* for an overview)
  - Modeling and prediction of electricity markets
- Here: Simplified approach
  - Description of average price and degree of variability
  - No predictive power

Idea: Categorize all prices into two regions, based on a threshold

High prices:
→ Temporary shutdown viable?

Low prices:
→ Normal operations

Price Model

$$k = \frac{p_{high}}{p_{avg}}$$

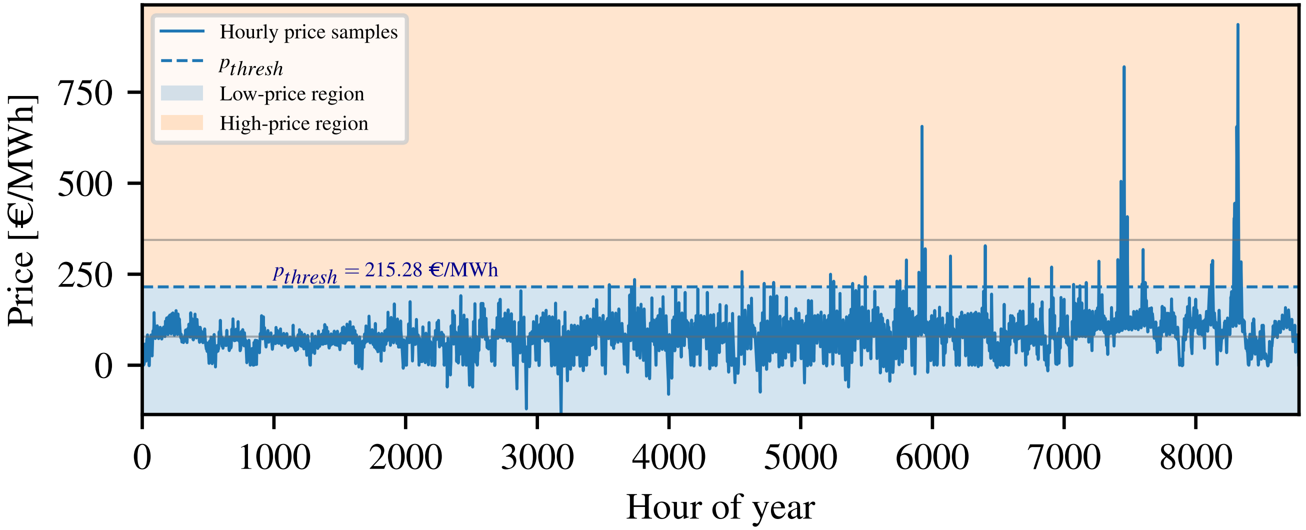

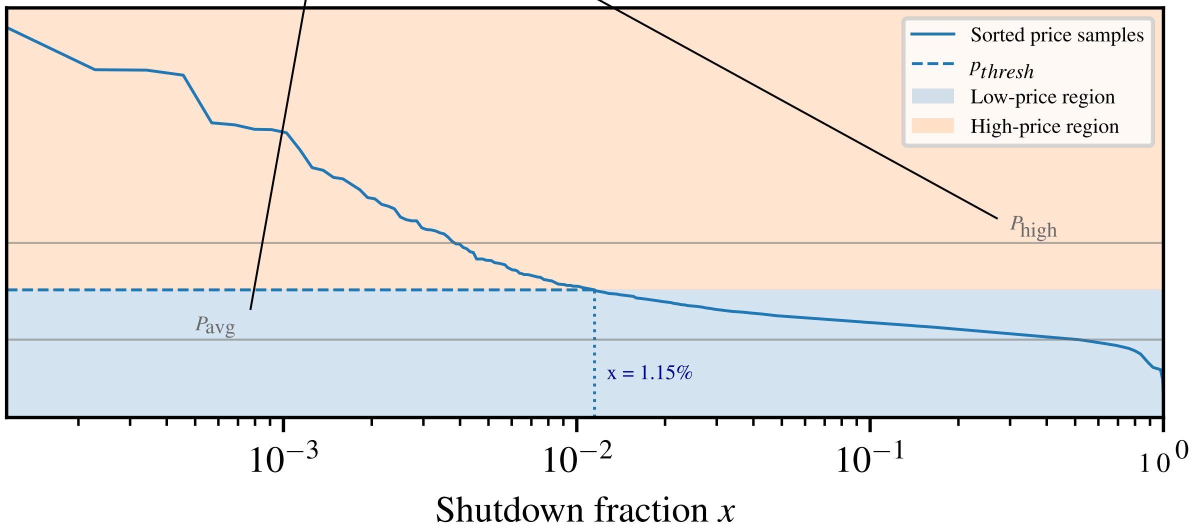

$$p_{high} = p_{avg} \cdot k$$

$$p_{low} = p_{avg} \cdot \frac{kx - 1}{x - 1}$$

Shutdown fraction x:
Fraction of time in which we consider prices to be high

Shutdown Policies

Two fundamental shutdown policies:

- *“Always on” (AO):* No shutdowns, regardless of energy price
- *“With shutdowns” (WS):* Complete shutdown when price is above threshold

Respective energy costs:

$$E_{AO} = T \cdot P \cdot p_{avg}$$

T: Overall time
P: Power consumption

$$\begin{aligned} E_{WS} &= T \cdot P \cdot (1-x) \cdot p_{low} \\ &= T \cdot P \cdot (1-x) \cdot p_{avg} \cdot \frac{kx - 1}{x-1} \\ &= T \cdot P \cdot p_{avg} \cdot (1-kx) \end{aligned}$$

Cost per Compute

- Comparison of fundamental shutdown policies:

  Quotient of the total costs and the time that the system is operational

  → *“Cost per compute”*

$$CPC_{AO} = \frac{F + E_{AO}}{T} \qquad CPC_{WS} = \frac{F + E_{WS}}{(1 - x) \cdot T}$$

Assessing Shutdown Viability

We can then express the question whether shutdowns are beneficial:

$$CPC_{WS} < CPC_{AO}$$

$$\frac{F + T \cdot P \cdot p_{avg} \cdot (1 - kx)}{(1 - x) \cdot T} < \frac{F + T \cdot P \cdot p_{avg}}{T}$$

$$k > \frac{F}{T \cdot P \cdot p_{avg}} + 1$$

$$k > \frac{F}{E_{AO}} + 1$$

Ratio $\frac{F}{E_{AO}}$: cost-distribution coefficient Ψ:

$$\Psi = \frac{F}{E_{AO}}$$

$$\boxed{k > \Psi + 1}$$

Case Study: Lichtenberg

Tier-2 HPC system at TU Darmstadt (~1,200 nodes)

- Estimate fixed costs (F)
  - Multiple generations simultaneously
  - Overlapping lifetimes

  → We consider average yearly investment
- Additionally: Long-term investments in building and cooling infrastructure
- Not considered: staff, etc.

- Estimate energy costs (E)
- Cost distribution coefficient: Fixed costs / Energy costs

$$\Psi_{LB} = \frac{F}{E_{AO}} \approx 2$$

Germany's spot-market electricity prices from 2024 on different time scales

Case study: Lichtenberg

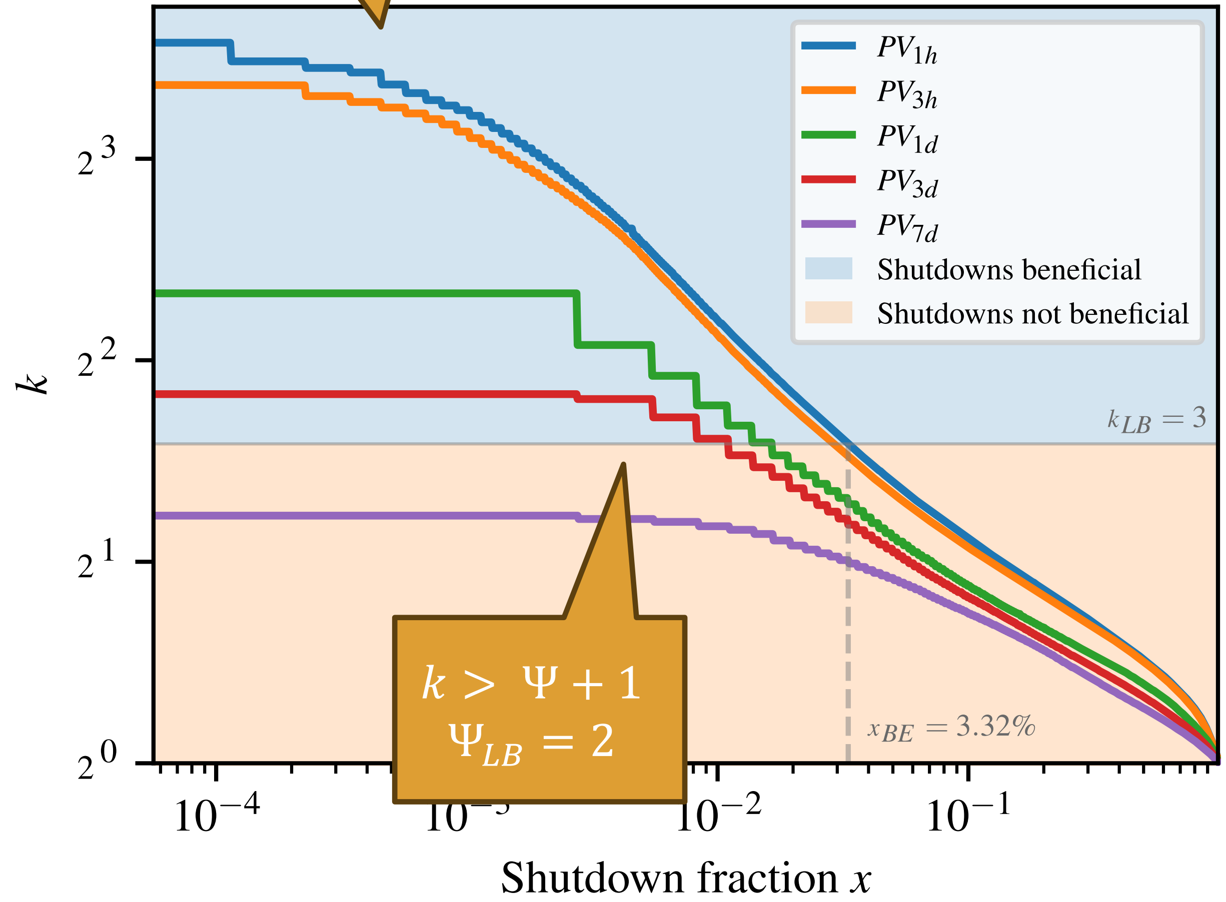

$k > \Psi + 1$
$\Psi_{LB} = 2$

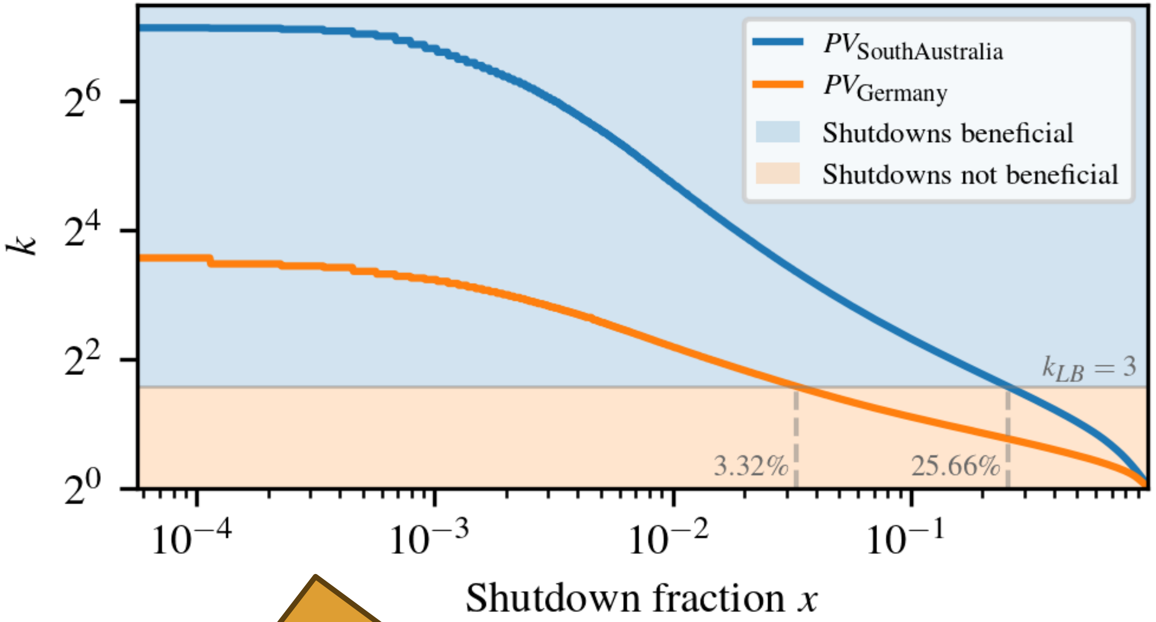

Comparison Germany vs. South Australia (time scale: 1 hour)

Case Study: Lichtenberg

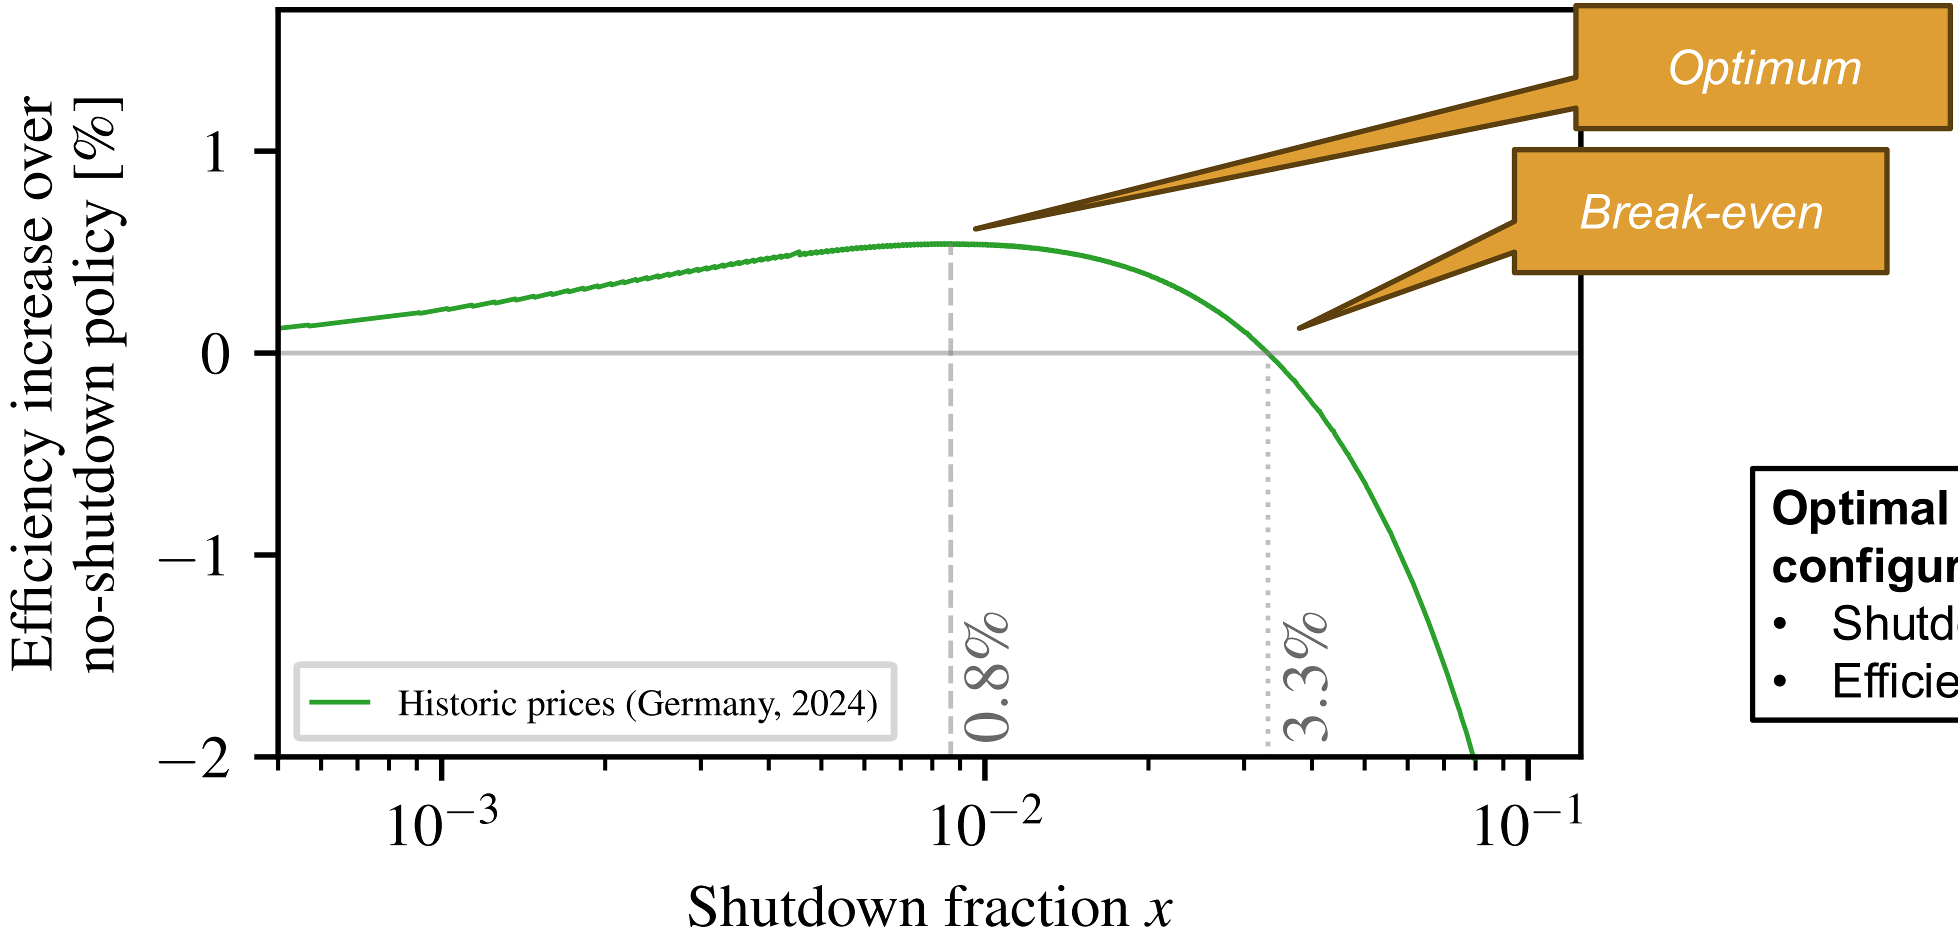

Optimal shutdown configuration:

- Shutdowns ~0.8% of time
- Efficiency increase: ~0.6%

Case Study: Variability

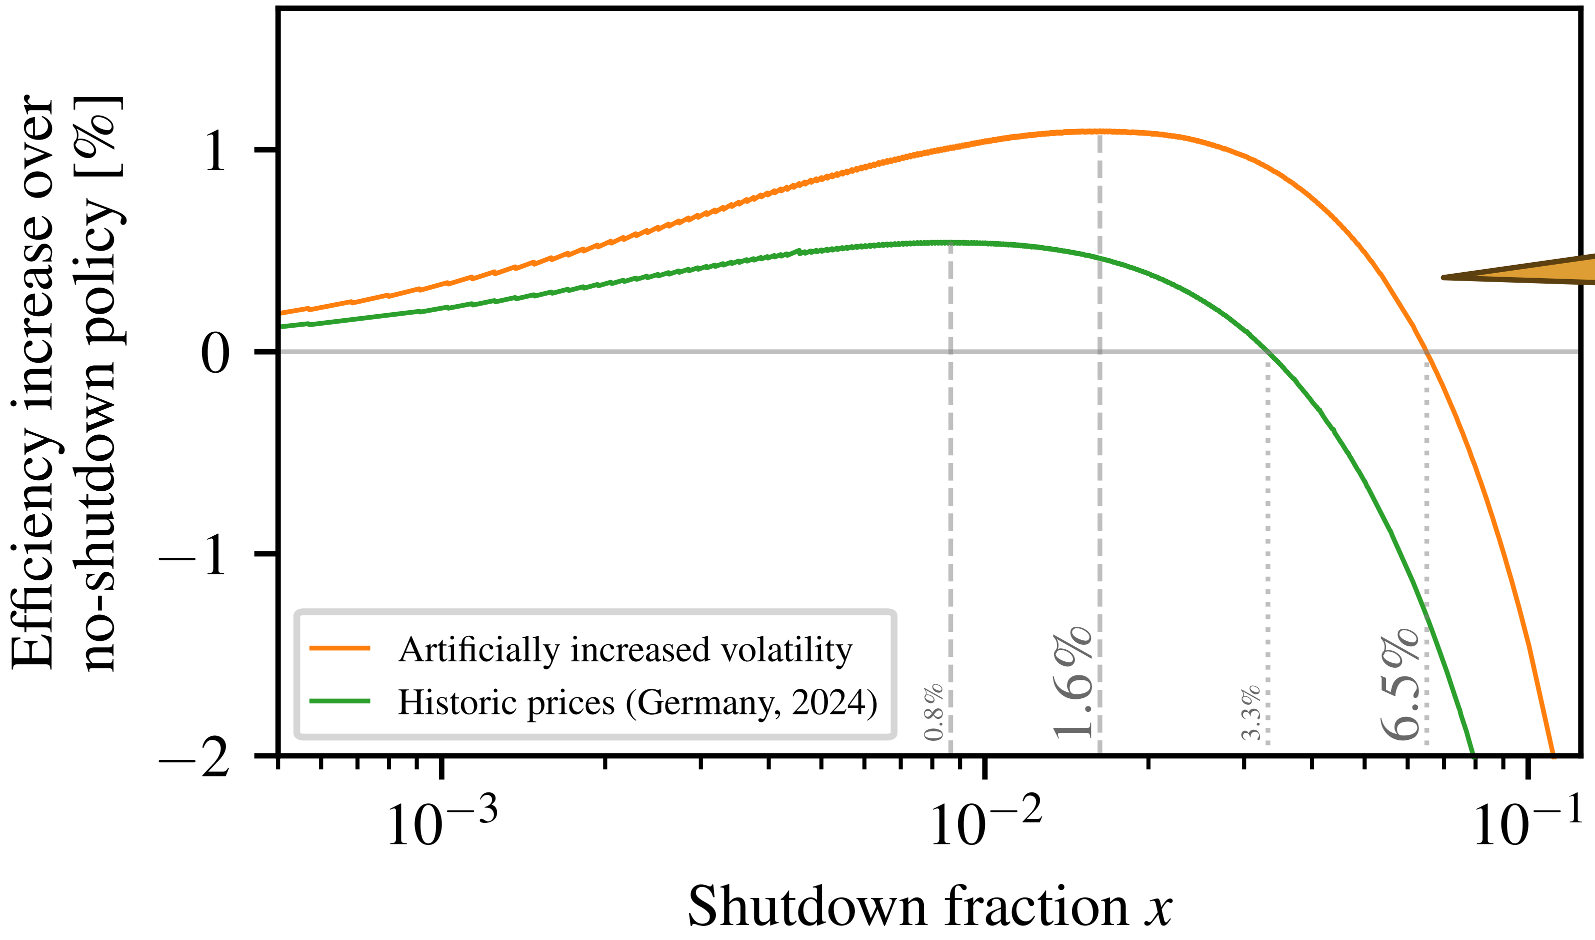

Hypothetical prices:
Renewable prices: -50%,
Fossil prices: +100%

Case Study: Cheaper Hardware

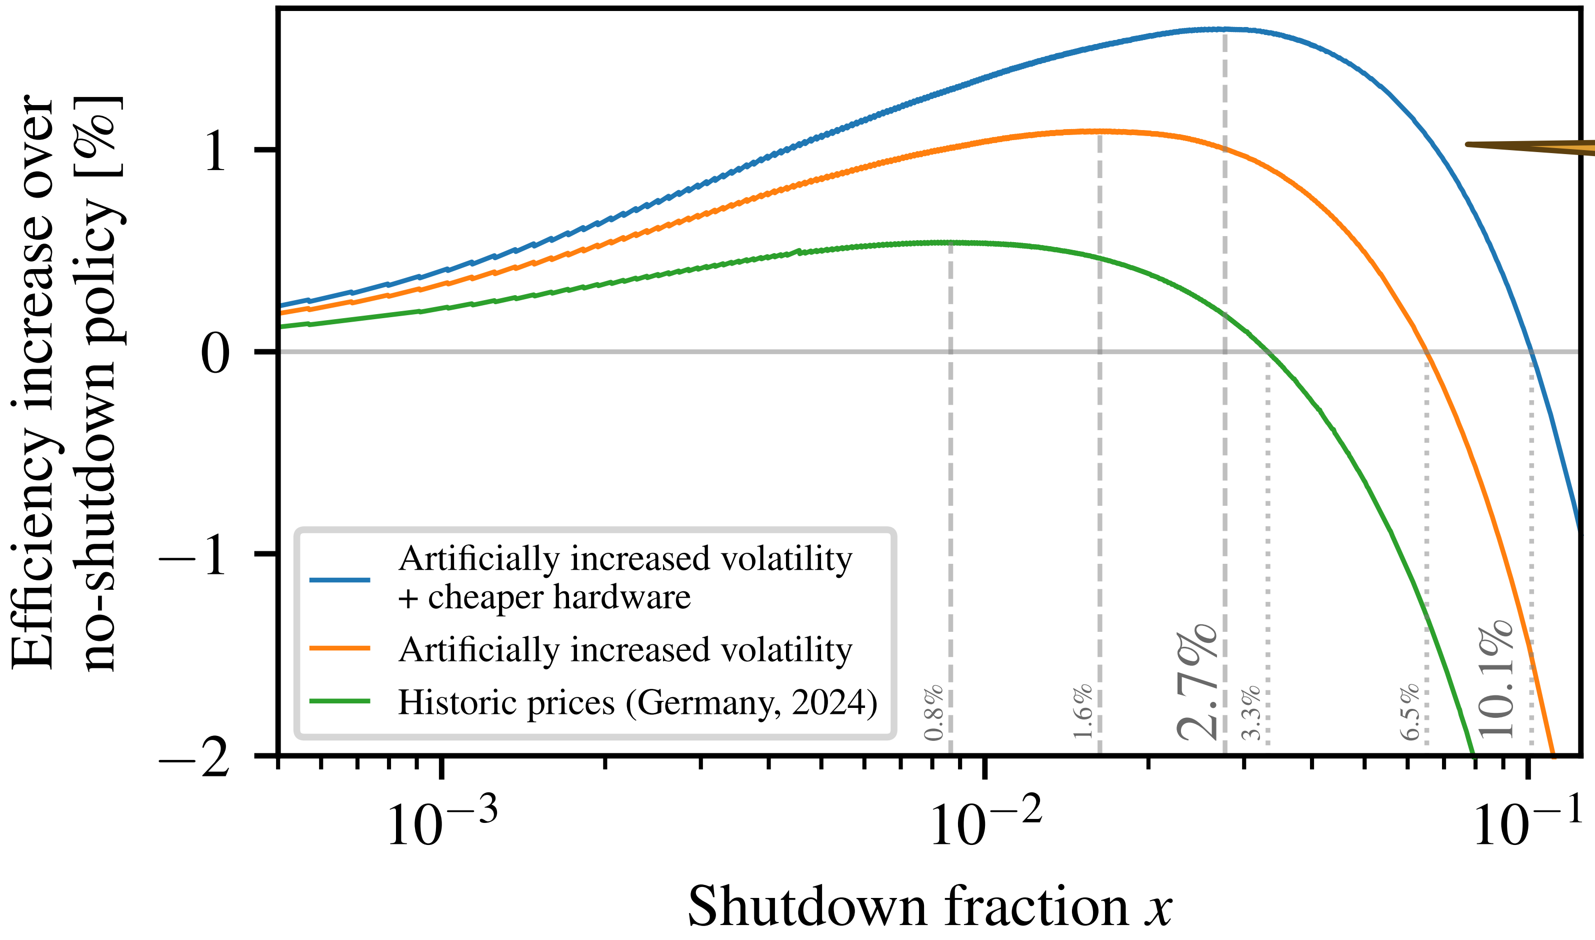

Ψ = 1.2 *instead of 2*

Case Study: Cheaper Hardware

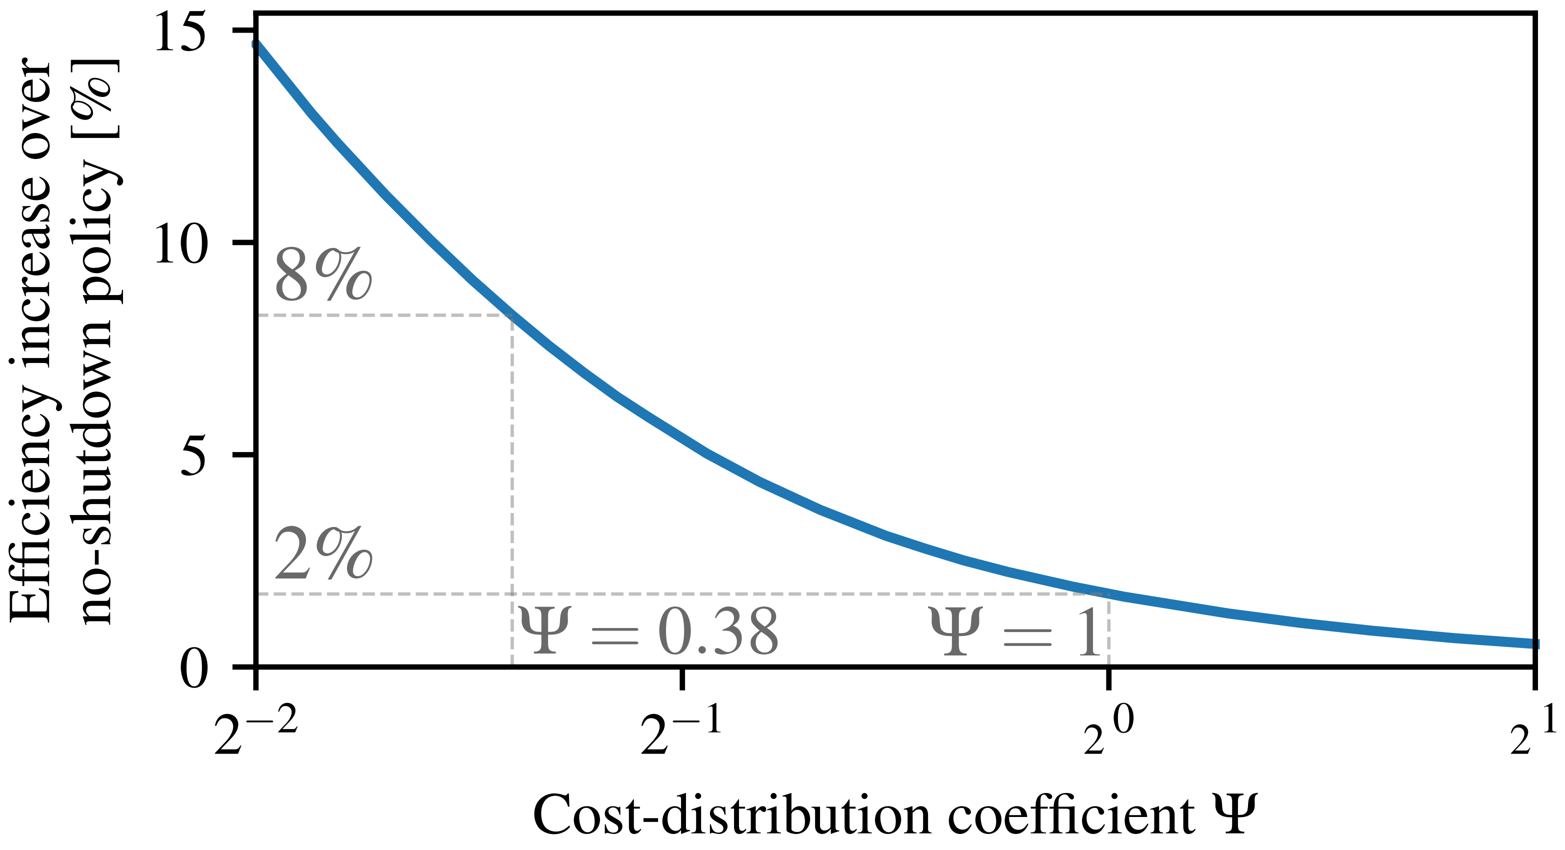

Case Study: Regional comparison

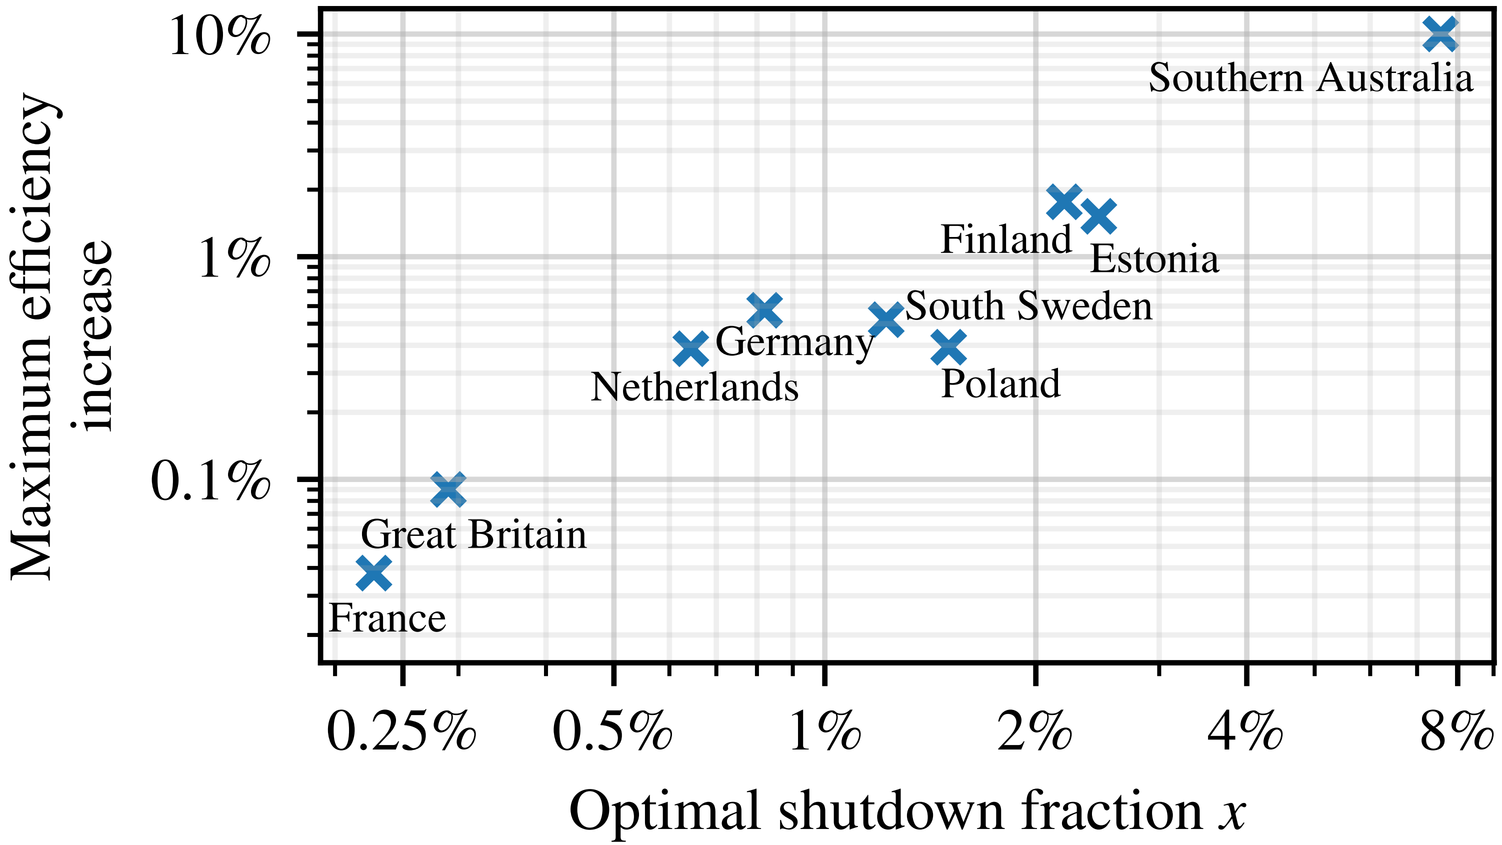

Discussion

- Assumption: Full hardware utilization when switched on

  → Necessary to isolate the effect of price variability

- Shutdown costs neglected

  → Upper bound of economic viability

- Optimization criterion and constraints
- Partial shutdowns

Discussion: Case Study Results

- For Lichtenberg: **only minimal efficiency gains in 2024, at best** (efficiency increase: < 0.54%)
- South Australian electricity market: Example for extreme price variability
- Future developments could strengthen economic case, either due to
  - Increased variability (e.g., higher CO2 taxes)
  - Shift in cost distribution

TECHNISCHE UNIVERSITÄT DARMSTADT

Parallel Programming

Summary

→ Model to estimate the economic viability of temporary shutdowns

- Assesses trade-off between energy savings and hardware underutilization
- Provides **upper bound for economic viability**

- Operational constraints introduce additional complexity

Outlook:

- Consideration of negative prices
- Technical and administrative hurdles

arXiv-preprint: Peter Arzt, Felix Wolf, *Navigating Energy Doldrums: Modeling the Impact of Energy Price Volatility on HPC Cost of Ownership*
https://arxiv.org/abs/2509.07567