

# From Resistive Load to Regulated Flexibility: Economic and Policy-Constrained Performance of AI-Based Mining-Heating Systems

**Javad VASHEGHANI FARAHANI**

Modul University Vienna, Austria

E-mail: [javad.vasheghani@modul.ac.at](mailto:javad.vasheghani@modul.ac.at)

*Received: 21 Feb. 2026 /Revised: 12 Mar. 2026 /Accepted: 16 Mar. 2026 /Published: 23 Mar. 2026*

**Abstract:** In terms of thermodynamics, proof-of-work blockchain mining is comparable to an electric resistive load whose electrical input is nearly completely dissipated as low-grade heat. This characteristic makes it possible to create mining–heating setups that serve two purposes: space heating and cryptographic computing. Nevertheless, rather than relying just on gross power use, the circumstances of the marginal grid, market exposure, and enforceable operational limits determine the environment and economic performance of such systems. In this study, mining-heating systems with AI are evaluated under certain legislative and economic constraints. A Bitmain Antminer S21 Pro (3.51 kW, 234 TH/s) is simulated using physics spanning realistic room volumes (60–340 m<sup>3</sup>) in a European comfort zone of 20–23 °C. The same thermal and economic assumptions are used to benchmark four control strategies: traditional electric resistance heating, hybrid modulation, reinforcement learning (Q-learning), and bang-bang. Price-aligned operation, marginal emissions conditions, and other governance-relevant situations are evaluated in addition to operational profit and comfort criteria. Bang-bang control yielded the largest comfort-valid profit under baseline deterministic settings (€108.41 at 160 m<sup>3</sup>), reinforcement learning (€103.73), hybrid modulation (€100.24), and electric resistance heating (–€116.77 at 60 m<sup>3</sup>). In its best comfort-valid scenario, reinforcement learning maintained 99.94 % time-in-band comfort and reduced duty to 95.69 % while achieving almost maximum profit. While bang-bang control generated the greatest best-case comfort-valid profit, hybrid control obtained the maximum comfort-feasible share over the entire simulated scenario set. Building-physics characteristics had lower secondary impacts, while network hashrate (–3.124 % profit per +1 % input), power price (–2.271 %), Bitcoin price (+2.115 %), and block reward (+2.115 %) dominated economic viability, according to the local elasticity analysis. Emissions results are conditional: mining-heated only lowers net system emissions when it replaces fossil fuel-based heating and is time-aligned with low-carbon or excess renewable power. The findings demonstrate that AI-controlled mining-heating systems may operate as programmable electric loads with potential flexibility value, but their feasibility is dependent on market alignment, marginal system conditions, and enforceable governance constraints. Reinforcement learning improved operational smoothness in some scenarios under deterministic training and assessment, but cross-seed robustness remained subpar, with 91.70 % ± 11.38 % time-in-band at 160 m<sup>3</sup> and 21 °C.

**Keywords:** Bitcoin mining, Dual-purpose heating, Demand-side flexibility, Marginal emissions, Energy market integration, Climate governance.

## 1. Introduction

The electrical intensity and accompanying emissions potential of proof-of-work (PoW) blockchain mining have made it a central topic in

discussions about energy and climate policy [1-3]. According to empirical estimates, Bitcoin's network-scale power consumption can approach that of small or medium-sized nation-states, with emissions results that are heavily influenced by the

underlying generating mix [1, 2]. Due to these worries, energy-intensive crypto-assets are now being assessed more and more in the context of larger sustainability and climate-governance frameworks in European legislative talks [4, 5].

For environmental evaluation, however, total electricity use is analytically inadequate on its own. Climate effect must be assessed at the margin, or based on how generating technology reacts to small variations in demand at a particular moment, in electrical systems that contain substantial percentages of variable renewable energy [2, 6].

In these circumstances, emissions results are determined by the temporal alignment of load with excess renewable power or fossil-based marginal supply. Therefore, it is crucial to distinguish between average and marginal generation impacts when evaluating the environmental performance of digital infrastructures that rely heavily on electricity [2, 6]. This marginal viewpoint is especially pertinent because of the features of bitcoin mining. Mining is technically interruptible, commercially price-sensitive, and geographically movable [1, 2]. The price of power, network hashrate, hardware efficiency, and block reward all have a direct impact on profitability [1, 3]. Within seconds, ASIC-based mining gear can reduce load from full capacity to almost nothing, with minimal physical damage and just lost income as a cost [7]. Full temporal relocation of demand is possible without disrupting downstream processes since mining production is uncertain and unrestricted by inventory or delivery dates [2]. These characteristics suggest that mining can operate as a highly flexible load or as a rigid baseload demand, dependent on governance design and market exposure.

A second aspect of analysis is related to thermodynamics. PoW computing is completely resistive at the system level as almost all electrical input dissipates as low-grade heat [2, 6]. Cooling systems are used in traditional data center layouts to reject this heat. However, when viewed through the lens of energy systems, such waste heat is recoverable thermal energy that is dependent on co-located demand. Mining gear can function as a computational co-product and an electric resistance warmer when included into building heating systems. The pertinent comparison in this dual-purpose setup is mining-with-heat-recovery vs the counterfactual heating technology and its marginal emissions profile rather than mining versus zero [2, 6].

European power markets have undergone structural changes that are impacted by this thermodynamic reinterpretation. Wind and solar capacity have grown rapidly, increasing the frequency of periods of surplus generation marked by curtailment and negative wholesale pricing. While congestion management expenses topped €4 billion in 2023, the European Union reduced over 12 TWh of renewable power [8, 9, 10]. 2023 saw a twelvefold rise in negative pricing occurrences, while 2024 statistics showed that Finland, Sweden, Germany, and the Netherlands had negative wholesale prices in about

8 % of hours, 7 % of hours, and 5 % of hours, respectively [8, 11]. Together with a lack of short-term demand-side flexibility, these patterns reflect temporal and geographical mismatches between renewable generation and demand [9, 12].

In these circumstances, flexible and price-responsive loads might provide value to the system by absorbing excess generation and withdrawing during times of shortage. However, flexibility only appears when there is proper market integration; it is not a feature of technology itself. Mining acts as stiff baseload demand and provides no systemic benefit in the absence of exposure to real-time pricing, curtailment signals, and enforced operational limits [2, 13].

Therefore, economic incentives and governance architecture, rather than technological identity, determine the impact of climate change. These factors interact with the decarbonization of space heating at the urban level. In the European Union, buildings consume around 40 % of all energy, with heating accounting for the majority of home energy usage [14]. Fossil fuels, especially natural gas, nevertheless contribute significantly to marginal heat provision in a number of Member States despite growing electrification [15]. Thus, dual-purpose mining-heating systems offer the possibility of combining local thermal service delivery with programmable electricity consumption.

Although proof-of-work mining's thermodynamic characteristics and potential for heat recovery are becoming more well acknowledged [2, 6], there is still a lack of thorough testing of building-integrated mining-heating systems under distinct control regimes and explicit market circumstances. Specifically, there hasn't been a comprehensive evaluation of the interplay of control strategy, economic sensitivity, and policy-relevant operational limitations using a single modeling approach.

Under the heading "Feasibility and Performance of Blockchain Dual-Purpose Mining-Heating Systems: Comparative Numerical Analysis of Bang-Bang, Hybrid, AI-Based, and Traditional Electric Heating Models," a preliminary version of this study was presented in the proceedings of the B2C'2025 conference. By adding policy-constrained evaluation, broader economic sensitivity analysis, and explicit integration with the characteristics of the European electricity market, the current study expands on earlier work [16].

In this work, a dual-purpose mining-heating system based on a Bitmain Antminer S21 Pro (3.51 kW, 234 TH/s) is simulated using physics. The same thermal and economic assumptions are used to assess conventional electric resistance heating, bang-bang control, hybrid modulation, and reinforcement-learning (Q-learning) control over realistic room sizes (60–340 m<sup>3</sup>) and a European comfort band of 20–23 °C. Duty factor, cycle frequency, time-in-band comfort, energy consumption, and profit are among the quantifiable operational parameters that are evaluated in this

research. In addition to technological viability, the model links operational performance to current market circumstances by taking into account sensitivity to the price of power, Bitcoin, network hashrate, and block reward. Additionally, findings are analyzed in the context of renewable surplus dynamics and marginal emissions in European power systems [8, 11], which enables evaluation of the circumstances in which mining-heating systems could operate as demand-side flexibility resources as opposed to fixed electrical loads.

This article conceptualizes mining-heating systems as programmable electricity demand whose climate and economic performance depends on market integration, marginal generation context, and enforceable operational constraints rather than on aggregate electricity consumption alone. It does this by combining thermodynamic conversion, AI-based control, economic sensitivity modeling, and governance-relevant system conditions within a single analytical framework.

## 2. Methodology

### 2.1. Modeling Approach and Scope

A deterministic reduced-order building energy model was created in order to assess the technical and financial viability of dual-purpose mining-heating systems. Under defined boundary conditions, the modeling technique combines proportionate proof-of-work revenue accounting, radiator-constrained heat transport, and thermodynamic heat balance.

In order to assess controller techniques while preserving physical interpretability, reduced-order single-zone models are frequently used in demand-response research and building energy modeling [7, 17]. In comparative controller analysis, where control and economic dynamics take precedence over spatial temperature gradients, such models are suitable.

In line with modern high-efficiency ASIC devices, the reference hardware is a Bitmain Antminer S21 Pro (234 TH/s; 3.51 kW rated input) [18]. Transient warm-up behavior, cycle frequency, steady-state functioning, and cumulative profit can all be analyzed thanks to simulations run at 1-minute resolution during 30-day winter periods (with the outside temperature set at 0 °C).

Simulations were carried out under deterministic boundary conditions, such as a set external temperature (0 °C), a constant power price, and a static network hashrate, in order to separate controller behavior and thermal dynamics. Without adding stochastic noise from weather unpredictability or market volatility, this deterministic setting allows for controlled comparison across controller systems. However, in real-world operations, network difficulty, electricity prices, and Bitcoin market value all change randomly. Therefore, rather than being predictive

forecasts of mining profitability, the current results should be viewed as structural baseline scenarios. The paradigm could be expanded in subsequent research by incorporating dynamic difficulty adjustment, weather variability, and stochastic power price series. In contrast, the current work uses local elasticity analysis around the baseline scenario to assess structural sensitivity.

European homes typically have room capacities between 60 and 340 m<sup>3</sup>. The EU building policy frameworks' comfort criteria are in line with the target interior temperatures of 20–23 °C [14].

### 2.2. Thermal Model Formulation

Indoor temperature evolution is governed by:

$$C \frac{dT}{dt} = Q_{in} - Q_{loss}, \quad (1)$$

where  $C$  is the effective thermal capacitance,  $Q_{in}$  is the delivered miner heat,  $Q_{loss}$  is a transmission and infiltration losses.

This lumped-capacitance formulation is consistent with the first-order equivalent thermal parameter (ETP) description of building thermal dynamics, which is frequently used in diagnostic analysis and control-oriented energy modeling [19].

Effective capacitance is defined as:

$$C = C_{air} \times V \times \varphi, \quad (2)$$

where  $C_{air} \approx 1200 \text{ J}/(\text{m}^3 \cdot \text{K})$ ,  $V$  is the room volume,  $\varphi = 3.0$ .

Using conventional thermophysical property data, the volumetric heat capacity of indoor air was estimated to be  $1200 \text{ J} \cdot \text{m}^{-3} \cdot \text{K}^{-1}$  [20]. An aggregate thermal-mass factor ( $\varphi$ ) was added to the lumped single-zone approximation to take into consideration extra thermal inertia from the building exterior and internal contents. Since reduced-order residential RC models often use values between 2 and 5,  $\varphi = 3.0$  was chosen as the mid-range value [21].

Envelope heat transfer coefficient [21]:

$$UA_{env} = UA + UA_{inf} \quad (3)$$

Envelope area is approximated assuming cubic geometry:

$$A = 6V^{2/3}, \quad (4)$$

In accordance with EU building type data and pre-renovation envelopes, the U-value is set at  $0.8 \text{ W}/\text{m}^2\text{K}$ , which represents the non-retrofitted European dwelling stock [14]. Infiltration losses are calculated using the standard air-change formula:

$$UA_{inf} = \dot{m}c_p, \quad (5)$$

$$\dot{m} = \rho \frac{ACHV}{3600} \quad (6)$$

where air density ( $\rho$  is represented by  $1.2 \text{ kg}\cdot\text{m}^{-3}$  and air changes per hour are indicated by ACH. This method is commonly used in thermal load estimates and building energy modeling, where infiltration heat loss is calculated by multiplying the mass flow rate by the specific heat of the air [21].

Electrical input at time  $t$  is defined as:

$$P_{elec,t} = f_t P_{rated} \quad (7)$$

Heat delivery is constrained by radiator transfer capacity:

$$Q_{cap} = UA_{rad} \max(T_{cap} - T_t, 0), \quad (8)$$

where  $T_{cap}$  represents the radiator's surface temperature and  $UA_{rad}$  the radiator's effective heat transfer coefficient.  $UA_{rad} = 120 \text{ W}\cdot\text{K}^{-1}$  was chosen as the value. This equates to around 3 kW of thermal output for a realistic temperature differential of 25 K. The standardized heat output ratings for residential panel radiators established by ISO 24365:2022 are in line with this magnitude [22].

In order to replicate low-temperature hydronic or forced-air heat recovery conditions appropriate for residential distribution systems, the radiator temperature cap was set to  $T_{cap} = 50 \text{ C}^\circ$ .

The room-coupling efficiency was set at  $\eta = 0.95$ , which accounts for the small distribution and convective losses that are common in thermally ducted or enclosed forced-air systems. Under regulated airflow regimes, integrated mining-thermal system modeling and experimental investigations show heat recovery efficiency of above 90 %.

In order to avoid unreasonably immediate thermal injection into the zone, this formulation enforces a physically restricted heat transfer rate [7].

Using proportionate network sharing logic, mining income is computed:

$$BTC/day = \left( \frac{H_{miner}}{H_{network}} \right) \times \bar{f} \times 144 \times B, \quad (9)$$

where  $H_{miner} = 234 \text{ TH}\cdot\text{s}^{-1}$ ,  $H_{network} \approx 915 \text{ EH}\cdot\text{s}^{-1}$ ,  $B = 3.125 \text{ BTC}$  (post-halving block subsidy), 144 represents the expected daily block production,  $\bar{f}$  denotes average operational duty fraction

BTC/day is multiplied by the observed BTC–EUR market price to get revenue in EUR [17, 23-25]. The cost of electricity is computed as follows:

$$C = E \times p_{elec} \quad (10)$$

Profit:

$$\Pi = R - C \quad (10)$$

Four control regimes assessed:

1. Electric resistance baseline;
2. Bang-bang on/off control;
3. Hybrid linear modulation;
4. Tabular Q-learning reinforcement learning.

A tabular Q-learning method was used to create the reinforcement learning controller. The action space included five discrete power fractions of the rated mining power:  $\{0, 0.25, 0.50, 0.75, 1.00\}$ . The indoor temperature was discretized into  $0.5 \text{ }^\circ\text{C}$  bins. The Q-table was conditioned on both the temperature state and the previously chosen action in order to deter fast switching behavior.

The reward function, which is defined as follows, gives operating smoothness and thermal comfort top priority:

$$r_t = r_c - \lambda_1 |T_t - T^*| - \lambda_2 (T_t - T^*)^2 - \lambda_3 \mathbb{1}(T_t > T^* + 2) - \lambda_4 |a_t - a_{t-1}|, \quad (12)$$

where  $T_t$  is the indoor temperature,  $T^*$  is the temperature setpoint, and  $a_t$  is the chosen power fraction. The reward coefficients were set to  $r_c = 2.0$  and  $\lambda_1 = 80.0$ ,  $\lambda_2 = 10.0$ ,  $\lambda_3 = 6000.0$ , and  $\lambda_4 = 2.0$ . This formulation discourages action modifications related to cycling and severely penalizes comfort breaches and overheating.

$\epsilon$ -greedy exploration decreased geometrically from 0.20 to 0.02 across 400 training episodes, with a learning rate of 0.12 and a discount factor of 0.995. The learnt policy was assessed under the same deterministic boundary circumstances on a 30-day horizon with 1-minute resolution after training on a 1-day horizon with 2-minute resolution.

Instead of focusing on direct profit maximization, the incentive formulation stresses thermal comfort and smooth actuation; steady operation and a lower duty factor led to indirect economic performance.

The RL controller's generalization was only evaluated in the predicted baseline scenario and not under stochastic price, weather, or network-difficulty dynamics because it was trained and evaluated under deterministic settings.

Symmetric  $\pm 10 \%$  parameter perturbations around the baseline example were used to assess local elasticities. This method is frequently used in energy-economic and techno-economic modeling to determine the main feasibility drivers since it allows for scale-independent assessment of relative effect across diverse factors.

Sensitivity tests were performed on two groups of parameters: Market determinants: Price of power, network hashrate, and Bitcoin Building determinants: envelope U-value, infiltration rate, and radiator heat transfer coefficient.

Due to the significant price and hashrate fluctuation that characterizes cryptocurrency mining economics [1, 2], elasticity research is required to differentiate between temporary market circumstances and structural viability. Such local sensitivity analysis offers a clear way to distinguish endogenous system

design factors from external market risk in the context of integrated energy systems [7, 12].

The results are generated under deterministic thermal and economic boundary conditions; dynamic network difficulty adjustments and stochastic power price fluctuation were not considered.

The reinforcement-learning controller was assessed using training diagnostics and cross-seed robustness statistics in addition to its final greedy policy in order to enhance repeatability and peer-review transparency. In particular, multi-seed assessment data for profit, energy, duty factor, cycle frequency, and time-in-band were collected, together with (i) the total training reward each episode and (ii) episode-wise action modifications as a surrogate for policy smoothness. By clearly separating best-case performance from cross-seed average behavior, these diagnostics were utilized to determine whether the learnt policy converged toward stable behavior under deterministic boundary conditions.

### 3. Results across Controllers

The technical and economic assumptions outlined in Section 2 (BTC €102701.28; power €0.0946/kWh; block reward 3.125 BTC; network hashrate 915 EH/s; pool fee 2 %) were used to all simulations. Assuring adherence to European thermal comfort requirements, results are presented for comfort-valid scenarios ( $\geq 95$  % of time between 20–23 °C).

Under the baseline deterministic assumptions, dual-purpose mining-heating systems produced positive operational profit in selected scenarios over the simulated room-volume range (60–340 m<sup>3</sup>) while preserving thermal comfort. The electric resistance baseline, on the other hand, produced no mining revenue and remained strictly cost-incurring despite being thermodynamically equivalent in conversion efficiency. This is in line with earlier energy-mining assessments that emphasized revenue dependence on hashrate share and market price [1, 6].

The results are given at two levels to prevent selection bias. Table 1 summarizes all 16 simulated volume-target combinations for each controller, regardless of comfort feasibility, and presents the share of cases that match the  $\geq 95$  % time-in-band requirement.

**Table 1.** Controller summary across all simulated scenarios (n = 16 per controller).

Controller	Mean Profit (€)	Mean Energy (kWh)	Mean Duty (%)
Bang-Bang Miner	93.98	2190.80	86.69
RL Miner (Q-learning)	86.82	2023.90	80.08
Hybrid Miner	92.96	2167.13	85.75
Electric Heater	-207.24	2190.73	86.69

Hybrid control had the largest comfort-feasible share (50.00 %) across all 16 scenarios per controller, followed by electric resistance heating (37.50 %), bang-bang control (31.25 %), and reinforcement learning (25.00 %).

**Table 2.** Best comfort-valid result per controller under baseline conditions.

Controller	Profit (€)	Energy (kWh)	Duty (%)
Bang-Bang Miner	108.41	2527.20	100.00
RL Miner (Q-learning)	103.73	2418.24	95.69
Hybrid Miner	100.24	2336.71	92.46
Electric Heater	-116.77	1234.31	48.84

With continuous full-duty operation, Bang-bang control generated the largest best-case comfort-valid profit (€108.41), but at the cost of the highest energy usage (2527.20 kWh). Reinforcement learning achieved near-maximal profit (€103.73), decreased duty to 95.69 %, maintained 99.94 % time-in-band comfort, and eliminated cycling in its best comfort-valid scenario. While significantly lowering energy consumption and duty factor, hybrid modulation produced a little lower best-case profit of €100.24. In its best-performing scenario, the electric-resistance baseline remained strictly costly while maintaining comfort.

Bang-bang control produced the fastest warm-up (4.0 min at 60 m<sup>3</sup>), indicating full-load operation. Depending on volume, reinforcement learning took 60–100 minutes to reach comfort, but it produced smoother trajectories and did away with cycling. Hybrid modulation decreased duty while suppressing cycling. Because frequent heat and electrical transients are linked to ASIC hardware deterioration, cycling frequency is operationally significant. By reducing mechanical and thermal stress, RL enhances quality-adjusted performance, whereas bang-bang optimizes immediate income.

Under the baseline assumptions, envelope heat-loss dominated at the largest simulated room volume (340 m<sup>3</sup>), requiring nearly constant operation to maintain comfort. A local elasticity study was performed around the 160 m<sup>3</sup> baseline instance in order to differentiate boundary-condition effects from structural factors. Elasticities were computed with  $\pm 10$  % symmetric perturbations of each parameter while maintaining the same values for the others. This method, which separates relative sensitivity, is frequently used in energy systems techno-economic assessments.

Economic viability is dominated by market factors. Profit is reduced by 3.124 % for every 1 % rise in network hashrate, which is more than the sensitivity to electricity prices (–2.271 %). In a similar vein, profit rises by 2.115 % for every 1 % increase in the price of Bitcoin or the block reward. These magnitudes support existing mining cost formulations [1, 26] by

confirming that power pricing and competitive network circumstances, rather than building physics factors, structurally drive mine–heating feasibility.

**Table 3.** Local profit elasticities around the RL baseline case (160 m<sup>3</sup>, 21 °C).

Parameter	Elasticity
Network hashrate	-3.124
Electricity price	-2.271
Bitcoin price	+2.115
Block reward	+2.115
Radiator UA	+0.604
Outdoor temperature	-0.581
Envelope U-value	+0.474
Infiltration ACH	+0.154

Secondary influence is exerted by building factors. Beyond around 110 W/K, radiator UA shows declining results, suggesting that heat-transfer limits are saturated. Near 0.9 W/m<sup>2</sup>K, the U-value exhibits a shallow non-monotonic response, indicating a relationship between operating duty and heat demand. Under constant electricity price, infiltration impacts are still minimal, indicating that envelope retrofits have a greater impact on comfort and energy efficiency than on mining profitability under static market assumptions.

**Table 4.** Multi-seed robustness of the RL controller at 160 m<sup>3</sup> and 21 °C (mean ± 95 % CI).

Metric	Mean ± 95 % CI
Profit (€)	103.19 ± 1.46
Energy (kWh)	2405.57 ± 34.08
Uptime (%)	99.54 ± 0.51
Duty (%)	95.19 ± 1.35
Cycles/day	3.99 ± 4.30
Time in band (%)	91.70 ± 11.38

The range of RL profit among seeds was €90.98 to €108.41, with a median of €103.73.

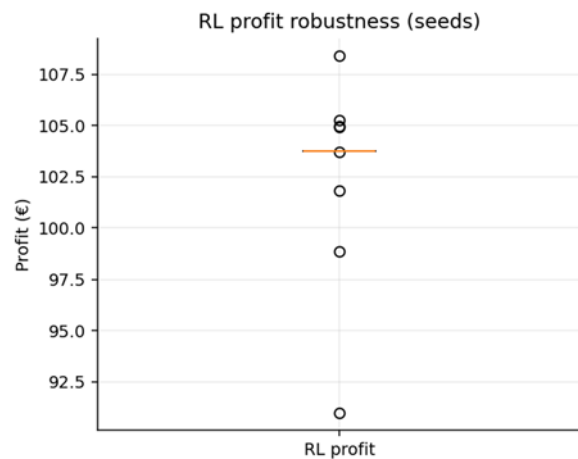
The cross-seed distribution of reinforcement-learning profit for the baseline scenario at 160 m<sup>3</sup> and 21 °C is shown in Fig. 1. The plotted points display the outcomes for individual random seeds, the middle line indicates the median, and the y-axis displays profit in euros. The figure supports Table 4 by showing that under deterministic boundary constraints, there is still discernible cross-seed variability even though the majority of runs cluster close to the median.

Reinforcement learning did not converge to the same results across random seeds, according to multi-seed robustness investigation at 160 m<sup>3</sup> and 21 °C. Rather, the controller attained a profit of €103.19 ± 1.46 (95 % CI), energy consumption of 2405.57 ± 34.08 kWh, uptime of 99.54 % ± 0.51 %, duty factor of 95.19 % ± 1.35 %, cycle frequency of 3.99 ± 4.30 cycles/day, and time-in-band of

91.70 % ± 11.38 %. Profit ranged from €90.98 to €108.41, with a median of €103.73. These findings show that although cross-seed comfort robustness is still lacking, reinforcement learning may identify high-performing strategies under deterministic boundary circumstances. Overall, the studies show three structural conclusions:

- Mining-heating systems can maintain comfort and produce positive operating profit under baseline economic assumptions in certain scenarios;
- Controller selection has a significant impact on operational quality, feasible operating envelope, and cycling behavior, even when best-case profits remain relatively close;
- Market variables are the primary determinants of economic viability, with building physics serving as a secondary modulation layer through thermal demand and comfort feasibility.

The mathematical basis for assessing mining-heating systems as programmable electric loads in larger energy and market settings is provided by these studies.



**Fig. 1.** RL profit robustness across seeds.

## 4. Discussion

### 4.1. Economic Structure and Revenue Exposure

According to the elasticity analysis, external market factors, not internal thermal characteristics, structurally dictate profitability. The price of power (-2.271 %) and network hashrate (-3.124 % per +1 %) have the most effects on profit, followed by the price of Bitcoin and the block reward (+2.115 %). This is consistent with well-known mining income formulas, where expenses scale linearly with energy consumption and predicted returns scale with relative hashrate share and block subsidy [1-3].

While they have an impact on duty factor and energy consumption, building-physics characteristics don't significantly change the revenue mechanism.

Rather from altering mining production, their influence is mediated via modulating heating demand. This demonstrates how dual-purpose mining-heating systems are still economically linked to wholesale energy prices and worldwide network rivalry. They are best described from an energy-economics standpoint as revenue-linked electric loads whose viability is dependent on the dynamics of competitive difficulty and the price ratios of power to Bitcoin [1, 3].

#### **4.2. Flexibility Value and Market Conditions**

The European energy markets are becoming more and more characterized by negative wholesale price intervals, congestion, and renewable unpredictability [10, 27]. Mostly in systems with strong wind and solar penetration, ACER expects a twelvefold rise in negative pricing occurrences in 2023, with further increases in 2024 [8]. These circumstances are a result of both a lack of short-term flexibility and fundamental temporal imbalances between generation and demand [8, 12].

Technically, mining hardware may quickly reduce load without causing process losses that go beyond lost revenue [2, 6]. Mining is distinct from traditional industrial demand, which is usually limited by process continuity needs, due to its high interruptibility. However, system value is not produced solely by technological prowess. Only when activities are subjected to real-time or short-interval pricing signals and react appropriately does flexibility become economically significant. Mining acts as stiff baseload demand in the absence of market exposure, which is consistent with more general results that market-access and regulatory restrictions restrict effective demand-side participation in Europe [8, 12, 13].

This means that, in reality, mining-heating systems would have to be directly exposed to intraday or day-ahead pricing, or aggregated into demand-response portfolios, in order to function as economically significant flexibility resources as opposed to passive electric heaters.

#### **4.3. Marginal Emissions and Counterfactual Assessment**

Instead of basing environmental assessment on total electricity usage, it must be done at the margin [2, 8]. The marginal generation technology's response at any given moment determines the emissions impact of incremental demand. While scarcity times are often determined by fossil generation, surplus intervals in renewable-rich systems may correlate to low or near-zero marginal emissions [6, 8].

Only when electricity use aligns with low-carbon or excess renewable generation and recovered heat

replaces fossil fuel-based heating can dual-purpose mining-heating lower net system emissions. Being close to renewable energy does not ensure a decrease in emissions. In accordance with network-scale emissions assessments of proof-of-work mining, the pertinent comparison is counterfactual displacement under particular temporal conditions [1, 6].

Since network hashrate and electricity prices have a significant impact on profitability, private economic incentives might not always coincide with emissions-optimal operation. Thus, it is essential to use policy tools that connect runtime to dynamic pricing, congestion signals, or flexibility market involvement in order to match system-level decarbonization goals with profitability [8, 12, 27].

Mandatory dynamic rates, automated curtailment compliance during system stress events, or eligibility requirements connected to supply contracts supported by renewable energy might all be used to operationalize this alignment.

#### **4.4. Urban Integration and Verifiable Operation**

About 40 % of the final energy consumption in the EU occurs in buildings, with space heating accounting for the majority of home demand [14, 15]. Thus, one of the main tenets of the European decarbonization policy is the electrification of heating. A programmable electricity-to-heat conversion process that is co-located with thermal demand is introduced by dual-purpose mining-heating.

It is most likely to be used in buildings taking part in aggregator-managed flexibility programs, gas-heated home retrofits, or electrification transition stages.

Operational findings demonstrate that, under some situations, it is possible to prevent excessive cycling and reduce the duty factor without sacrificing comfort. Reinforcement learning reached almost maximum profitability with 0 % cycling in its most comfortable-valid scenario, indicating that digital control may influence both thermal service quality and electrical-load smoothness. However, this finding should be understood as conditional rather than universally resilient across training initializations because the multi-seed study showed non-negligible heterogeneity in comfort performance. Therefore, in addition to enhanced control, market exposure and enforceable operational restrictions are necessary for the creation of flexibility value. AI-controlled mining-heating systems provide programmable electric loads with conditional flexibility potential, especially when integrated into aggregator-based demand-response frameworks or price-responsive household electrification.

#### **4.5. Structural Constraints**

Transmission expansion, storage deployment, or fundamental market change cannot be replaced by

programmable digital demand. Large-scale renewable integration requires not only responsive end-use demand but also wider flexibility resources and grid reinforcement, according to system assessments conducted throughout Europe. Although they may respond to price signals or lessen short-term restriction, flexible mining-heating loads are nevertheless constrained by broader infrastructure and market factors.

Furthermore, profitability is still largely dependent on the state of the worldwide network. Regardless of local heating demand or building performance, revenue share decreases as network hashrate increases. This external reliance restricts long-term economic stability and sets mining-heating apart from traditional electrification methods. Therefore, rather than being intrinsically sustainable or system-beneficial technologies, dual-purpose mining-heating systems should be seen as conditional flexibility mechanisms.

## 5. Conclusion

This study used a physics-based reduced-order thermal model in conjunction with proportional proof-of-work revenue accounting to assess the economic, operational, and policy-constrained performance of dual-purpose mining-heating systems. Certain mining-heating configurations were able to maintain the 20–23 °C comfort band while producing positive operating profit across representative residential room volumes under baseline deterministic conditions (BTC €102701.28; electricity €0.0946/kWh; block reward 3.125 BTC; network hashrate 915 EH/s; pool fee 2 %). In contrast, because the electric-resistance baseline offered thermal service without mining revenue, it continued to be strictly cost-incurring.

Operational quality was more significantly impacted by controller choice than structural profitability. Bang-bang control produced the greatest profit (€108.41) among the best comfort-valid scenarios, but it also used the most energy and had the least operational moderation. In the best-performing comfort-valid instance, reinforcement learning eliminated cycling, lowered duty factor to 95.69 %, maintained 99.94 % time-in-band comfort, and earned nearly maximum best-case profit (€103.73). In all of the tested scenarios, hybrid modulation obtained the biggest comfort-feasible share and earned a little lower best-case profit (€100.24) while significantly reducing energy use and duty factor. These results suggest that rather than significantly altering the income structure, control design largely affects cycling behavior, smoothness, and the practical operating range.

The sensitivity analysis demonstrated that market factors are the primary determinants of economic viability. The price of power and network hashrate had the biggest detrimental influence on profit, whereas the price of Bitcoin and block reward had similarly significant beneficial effects. The lesser secondary

effects of building-physics parameters, such as radiator UA, external temperature, envelope U-value, and infiltration rate, were mostly driven by thermal demand and comfort feasibility. Therefore, rather than being solely dependent on building performance, dual-purpose mining-heating is nevertheless fundamentally linked to wholesale electricity pricing and international mining rivalry.

Conditional rather than intrinsic flexibility value determines these systems' policy relevance. Only when recovered heat replaces fossil-based heating and energy demand is temporally matched with low-carbon or excess renewable power can mining-heating minimize net system emissions. Similarly, only when an operation is subject to curtailment signals, dynamic pricing, or other enforceable market-based control frameworks do flexibility benefits become apparent. Without these restrictions, mining acts more like a revenue-seeking electric load than a resource with guaranteed flexibility.

Additionally, any strong assertion of resilience is qualified by the reinforcement-learning outcomes. With time-in-band equal to  $91.70 \% \pm 11.38 \%$  at 160 m<sup>3</sup> and 21 °C, multi-seed evaluation revealed non-negligible variability in comfort performance, despite the RL controller identifying high-performing policies in some deterministic circumstances. This suggests that while RL-based control might enhance operational smoothness under certain assumptions, its dependability is still insufficient to substantiate claims of generally steady performance.

In general, dual-purpose mining-heating systems should be understood as conditional programmable electric loads whose environmental and economic performance is contingent upon market exposure, marginal grid conditions, controller design, and enforced operating limits. Their most likely function is not as intrinsically sustainable heating technologies, but rather as specialized flexibility-linked systems that could prove beneficial in carefully regulated building integration and the electricity market. In order to evaluate policy and operational claims under more realistic system settings, future work should expand the current deterministic framework by including stochastic power prices, dynamic network difficulty, weather variability, and explicit marginal-emissions time series.

## Acknowledgements

My profound appreciation goes out to Prof. Dr. Horst Treiblmaier for his invaluable supervision, direction, and unwavering support during this research. His wise counsel, support, and constructive criticism have been crucial to the accomplishment of this project. Funding: This research was funded in whole by the Austrian Science Fund (FWF), The Energy Use of Bitcoin, Grant DOI: 10.55776/PAT9707423.

## References

- [1]. M. A. Rudd, R. M. Moore, Bitcoin and its energy, environmental, and social impacts: An assessment of key research needs in the mining sector, *Challenges*, Vol. 14, Issue 4, 2023, 47.
- [2]. H. Treiblmaier, A comprehensive research framework for Bitcoin's energy use: Fundamentals, economic rationale, and a pinch of thermodynamics, *Blockchain: Research and Applications*, Vol. 4, Issue 3, 2023, 100149.
- [3]. J. Li, N. Li, J. Peng, H. Cui, et al., Energy consumption of cryptocurrency mining: A study of electricity consumption in mining cryptocurrencies, *Energy*, Vol. 168, 2019, pp. 160-168.
- [4]. European Parliament, Cryptocurrencies in the EU: new rules to boost benefits and curb threats, 2022, <https://www.europarl.europa.eu/news/en/press-room/20220309IPR25162/cryptocurrencies-in-the-eu-new-rules-to-boost-benefits-and-curb-threats>
- [5]. Finansinspektionen, Crypto-assets are a threat to the climate transition – energy-intensive mining should be banned, 2021, <https://www.fi.se/en/published/presentations/2021/crypto-assets-are-a-threat-to-the-climate-transition--energy-intensive-mining-should-be-banned/>
- [6]. A. Krause, T. Tolaymat, Quantification of energy and carbon costs for mining cryptocurrencies, *Nature Sustainability*, Vol. 1, 2018, pp. 711-718.
- [7]. W.-H. Chen, F. You, Energy optimization of bitcoin mining integrated greenhouse with model predictive control, *Applied Energy*, Vol. 395, 2025, 126256.
- [8]. ACER, Key developments in European electricity and gas markets 2025 Monitoring Report, 2025, [https://www.acer.europa.eu/sites/default/files/documents/Publications/2025\\_ACER\\_Gas\\_Electricity\\_Key\\_Developments.pdf](https://www.acer.europa.eu/sites/default/files/documents/Publications/2025_ACER_Gas_Electricity_Key_Developments.pdf)
- [9]. J. Gorenstein Dedecca, M. Ansarin, C. Bene, T. Van Delzen, et al., Increasing flexibility in the EU energy system, Policy Department for Transformation, Innovation and Health, European Parliament, 2025, [https://www.europarl.europa.eu/RegData/etudes/STUD/2025/769347/ECTI\\_STU\(2025\)769347\\_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/STUD/2025/769347/ECTI_STU(2025)769347_EN.pdf)
- [10]. OECD, Grids: Diagnostic toolkit for reducing regulatory barriers to solar, wind and pumped hydro storage in the European Union, [https://www.oecd.org/en/publications/diagnostic-toolkit-for-reducing-regulatory-barriers-to-solar-wind-and-pumped-hydro-storage-in-the-european-union\\_15f4aed4-en/full-report/grids\\_e2c2091b.html](https://www.oecd.org/en/publications/diagnostic-toolkit-for-reducing-regulatory-barriers-to-solar-wind-and-pumped-hydro-storage-in-the-european-union_15f4aed4-en/full-report/grids_e2c2091b.html)
- [11]. International Energy Agency, Electricity 2025 – Prices, 2025, <https://www.iea.org/reports/electricity-2025/prices>
- [12]. A. Mellot, M. Pollitt, I. Staffell, Electrification, flexibility or both? Emerging trends in European energy policy, *Energy Policy*, Vol. 206, 2025, 114725.
- [13]. Market monitor for demand side flexibility, *smartEn, LCP Delta*, 2025.
- [14]. European Commission, Energy Performance of Buildings Directive, [https://energy.ec.europa.eu/topics/energy-efficiency/energy-performance-buildings/energy-performance-buildings-directive\\_en](https://energy.ec.europa.eu/topics/energy-efficiency/energy-performance-buildings/energy-performance-buildings-directive_en)
- [15]. Eurostat, Renewable energy statistics, [https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Renewable\\_energy\\_statistics](https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Renewable_energy_statistics)
- [16]. J. Vasheghani Farahani, Feasibility and performance of blockchain dual-purpose mining–heating systems, in *Proceedings of the 4<sup>th</sup> Blockchain and Cryptocurrency Conference (B2C'25)*, 2025, pp. 13-16.
- [17]. Austrian Power Grid, Day-Ahead prices, <https://markt.apg.at/en/transparency/cross-border-exchange/day-ahead-prices/>
- [18]. BITMAIN Technologies, S21 Pro specification, <https://support.bitmain.com/hc/en-us/articles/31321354157593-S21-Pro-Specification>
- [19]. R. Sonderegger, Diagnostic tests determining the thermal response of a house, Technical Report LBL-7815, *Lawrence Berkeley National Laboratory*, 1978.
- [20]. P. Michalak, Impact of air density variation on a simulated earth-to-air heat exchanger's performance, *Energies*, Vol. 15, Issue 9, 2022, 3215.
- [21]. G. Fraisse, C. Viardot, O. Lafabrie, G. Achard, Development of a simplified and accurate building model based on electrical analogy, *Energy and Buildings*, Vol. 34, Issue 10, 2002, pp. 1017-1031.
- [22]. ISO 24365:2022 – Radiators and convectors – Methods and rating for determining the heat output, *International Organization for Standardization*, 2022.
- [23]. Blockchain.com, Bitcoin total hash rate (TH/s), <https://www.blockchain.com/explorer/charts/hash-rate>
- [24]. Hashrate.no, WhitePool mining pool for cryptocurrencies, <https://www.hashrate.no/pools/WhitePool>
- [25]. Yahoo Finance, BTC–EUR historical price data, <https://de.finance.yahoo.com/quote/BTC-EUR/>
- [26]. I. Staffell, S. Pfenninger, N. Johnson, A global model of hourly space heating and cooling demand at multiple spatial scales, *Nature Energy*, Vol. 8, Issue 12, 2023, pp. 1328-1344.
- [27]. European Commission, Reform of the EU electricity market design, 2023, [https://ec.europa.eu/commission/presscorner/detail/en/ip\\_23\\_1591](https://ec.europa.eu/commission/presscorner/detail/en/ip_23_1591)

