LUISS

Department of Economics and Finance

Chair of Sustainable Finance

Nuclear Energy as Future Infrastructure for Artificial Intelligence: Evidence from U.S. Electric Utilities

Supervisor

Prof. Dominik Damast

Co-Supervisor

Prof. Claudio Boido

Candidate

Gabriele Kamran Paschos Gricia

783021

Academic Year 2025/2025

Abstract

The rapid expansion of Artificial Intelligence (AI) has intensified concerns regarding the long-term electricity requirements of large-scale data centers and computational infrastructure. Nuclear energy has re-emerged in policy and industry debates as a potential stable and low-carbon solution capable of supporting AI-driven energy demand. This thesis investigates whether capital markets already price nuclear energy exposure as a strategically relevant asset characteristic in the context of AI industry development.

Using a sample of U.S. electric utilities from January 2016 to January 2026, firms are classified into nuclear-exposed and non-nuclear groups based on their core generation technology. Equal-weighted portfolios are constructed, and a Nuclear Long–Short (NLS) factor is derived to capture relative performance. Risk-adjusted returns are evaluated using the Fama–French three-factor model. In addition, public attention is proxied by Google Search Volume Index (SVI) data to examine whether nuclear valuation is state-dependent.

The results show that nuclear-exposed firms do not generate persistent abnormal returns relative to non-nuclear utilities over the full sample period. Differences in performance are largely explained by systematic factor exposures. However, nuclear firms significantly outperform during periods of elevated public attention, and pooled regressions confirm that attention interacts positively with nuclear exposure in explaining excess returns.

These findings suggest that nuclear energy is not yet priced as a structural AI infrastructure premium. Instead, its valuation appears conditional and narrative-sensitive, fluctuating with macro-attention cycles rather than reflecting a permanent repricing linked to AI-driven electricity demand.

Contents

Abstract **ii**

1 Introduction **1**

1.1 The Energy Constraint of Artificial Intelligence . . . 1

1.2 The Strategic Relevance of Nuclear Energy in the AI Era . . . 2

1.3 Research Objectives and Research Questions . . . 3

1.4 Structure of the Thesis . . . 5

1.5 Final Considerations . . . 5

2 Literature Review **7**

2.1 The Compute–Energy Constraint of Artificial Intelligence . . . 7

2.2 Energy Transitions and the Role of Firm Low-Carbon Capacity . . . 8

2.3 Nuclear Energy as Capital-Intensive Infrastructure . . . 9

2.4 Asset Pricing of Utilities and Infrastructure Firms . . . 10

2.5 Investor Attention, Expectations, and Google Trends . . . 11

2.6 Synthesis and Implications for the Empirical Strategy . . . 11

3 Data and Methodology **13**

3.1 Firm Classification . . . 13

3.1.1 Nuclear-Exposed Firms . . . 13

3.1.2 Non-Nuclear Firms . . . 13

3.2 Equity Data and Return Construction . . . 14

3.2.1 Price Data . . . 14

3.2.2 Return Computation . . . 14

3.3 Portfolio Construction . . . 14

3.3.1 Equal-Weighted Nuclear Portfolio 14
3.3.2 Equal-Weighted Non-Nuclear Portfolio 15
3.3.3 Long–Short Factor Construction 15
3.4 Fama–French Three-Factor Regression Framework 15
3.4.1 Factor Data 15
3.4.2 Regression Specifications 15
3.4.3 Estimation Procedure 16
3.5 Performance Metrics 16
3.6 Google Trends and Attention-Based Analysis 17
3.6.1 Google Trends Data 17
3.6.2 Attention Regimes 17
3.6.3 High–Low Return Differences 17
3.6.4 Pooled Regression 17
3.7 Methodological Considerations 18

4 Empirical Results **19**
4.1 Descriptive Performance of Portfolios 19
4.1.1 Return and Volatility 19
4.1.2 Cumulative Returns 20
4.1.3 Drawdown Analysis 21
4.1.4 NLS Cumulative Performance and Drawdowns 22
4.2 Fama–French Three-Factor Results 24
4.2.1 Regression estimates 24
4.2.2 Interpretation of factor exposures 24
4.3 Attention Regimes and Google Trends Evidence 25
4.3.1 High versus Low Attention Regimes 26
4.3.2 Pooled Regression with Attention Interaction 28

5 Conclusion **30**
5.1 Summary of Findings and Answer to the Research Question 30
5.2 Interpretation and Broader Implications 31
5.3 Limitations and Future Research 32

A Classification of Nuclear and Non-Nuclear Firms **36**

A.1 Nuclear-Exposed Firms 37

A.2 Non-Nuclear Firms 38

List of Figures

4.1 Cumulative performance (base 1) of NuclearEW and NonNuclearEW, January 2016 to January 2026. 21
4.2 Drawdowns of NuclearEW and NonNuclearEW, January 2016 to January 2026. 22
4.3 Cumulative performance (base 1) of NuclearEW, NonNuclearEW, and NLS, January 2016 to January 2026. 23
4.4 Drawdowns of NuclearEW, NonNuclearEW, and NLS, January 2016 to January 2026. 23
4.5 Mean monthly returns across High and Low attention regimes. 26
4.6 Annualized mean returns across High and Low attention regimes. 27

List of Tables

4.1 Descriptive Performance Statistics 20

4.2 Fama–French Three-Factor Regression Results 24

4.3 Mean Monthly Returns by Attention Regime 26

4.4 Annualized Mean Returns by Attention Regime 27

4.5 Pooled OLS with Attention Interaction 28

A.1 Nuclear-Exposed Tickers and Company Names 37

A.2 Non-Nuclear Tickers and Company Names 38

1 Introduction

1.1 The Energy Constraint of Artificial Intelligence

Over the last decade, artificial intelligence has evolved from a specialized research field into large scale industrial deployment. A defining characteristic of this transition is the rapid increase in computational requirements used to train frontier models. Amodei and Hernandez document that the amount of compute used in the largest AI training runs has grown at an annualized rate that far exceeds the historical pace of hardware improvements traditionally associated with Moore's Law (Amodei and Hernandez, 2018). Subsequent research on scaling laws formalizes this relationship, showing that model performance improves predictably as a function of increases in parameters, training data, and computational resources (Kaplan et al., 2020; Hoffmann et al., 2022). These findings imply that advanced AI development is directly constrained by the availability of computational infrastructure. Hardware level innovation has supported this expansion through the introduction of domain specific accelerators optimized for deep learning workloads, including tensor processing units and specialized supercomputing architectures (Jouppi et al., 2018; Jouppi et al., 2020). However, improvements in chip efficiency do not eliminate the underlying dependence on electricity. Instead, they enable larger models to be trained, reinforcing total energy demand. The physical manifestation of this growth is the rapid expansion of hyperscale data centres. AI training clusters operate at high utilization rates and require continuous and stable power delivery. According to the International Energy Agency, electricity consumption from data centres is projected to increase significantly over the coming decade, with artificial intelligence identified as a major driver of incremental demand (IEA, 2024). Private sector projections similarly anticipate substantial increases in power requirements associated with AI oriented infrastructure

(Goldman Sachs, 2024). The relevance of this trend lies in its structural character. Empirical evidence on scaling laws indicates that performance gains remain strongly correlated with increases in compute (Kaplan et al., 2020; Hoffmann et al., 2022). If this relationship persists, future improvements in artificial intelligence will require sustained expansion of electricity intensive infrastructure. Under this scenario, energy availability becomes a foundational input to technological progress rather than a secondary constraint. Artificial intelligence should therefore be understood not only as a software innovation but as an infrastructure dependent technology. As electricity demand from AI continues to rise, the characteristics of the energy system, including reliability, scalability, and cost structure, may influence the trajectory of AI driven growth. The next section examines whether nuclear energy, as a source of continuous low carbon baseload power, occupies a strategic position within this evolving framework.

1.2 The Strategic Relevance of Nuclear Energy in the AI Era

The projected expansion of artificial intelligence driven electricity demand raises the question of which energy technologies can reliably sustain large scale and continuous power consumption. Electricity systems differ significantly in their ability to provide stable output, manage variability, and support high load factors. For infrastructure intensive applications such as AI data centres, reliability and continuity of supply are central operational requirements. Renewable energy sources such as solar and wind have experienced rapid deployment and declining costs over the past decade. However, their output is inherently variable and dependent on weather conditions. The economic literature documents that increasing shares of variable renewables in the electricity mix can reduce their market value due to intermittency and correlation effects (Hirth, 2013). System level analyses further show that high penetration of wind and solar requires complementary flexibility resources, including storage, transmission expansion, or firm generation capacity (Reichenberg et al., 2018; Sepulveda et al., 2018). Although battery storage technologies have improved and may play an important balancing role, their large scale deployment involves significant cost considerations and system integration constraints

(Zakeri and Syri, 2015; Mallapragada et al., 2020). Nuclear energy differs structurally from variable renewable generation. Nuclear power plants operate with high capacity factors and provide continuous electricity output independent of climatic conditions. According to comparative cost assessments of electricity generation technologies, nuclear power remains one of the primary sources of firm low carbon electricity in advanced economies (IEA and OECD NEA, 2020; Lazard, 2025). From a system perspective, firm generation resources reduce the need for large scale balancing infrastructure and contribute to grid stability in high electrification scenarios (Sepulveda et al., 2018). Recent institutional reports have highlighted the potential compatibility between advanced nuclear technologies and data centre electricity demand. The International Atomic Energy Agency and the U.S. Department of Energy both note growing interest in pairing nuclear generation with large scale digital infrastructure, particularly where continuous power supply and carbon constraints coexist (IAEA, 2025; DOE, 2025). While these discussions remain partly prospective, they reflect a broader recognition that AI expansion may require firm and scalable electricity sources. The financial relevance of these technological characteristics lies in investor expectations. If artificial intelligence generates structurally higher electricity demand and if nuclear energy is perceived as a reliable provider of baseload power, then utilities with nuclear exposure may benefit from differential valuation dynamics. Capital markets may anticipate stronger long term demand, improved asset utilization, or strategic positioning within an AI intensive economy. This thesis does not assume that nuclear energy is the only viable solution for supporting artificial intelligence. Instead, it investigates whether financial markets price nuclear exposure as a strategic attribute under conditions of rising AI attention and projected electricity demand growth. By focusing on publicly traded utilities, the analysis examines whether differences in energy generation portfolios are reflected in risk adjusted return patterns.

1.3 Research Objectives and Research Questions

The previous sections establish two core premises. First, artificial intelligence development is increasingly dependent on large scale computational infrastructure, which in turn requires substantial and continuous electricity supply. Second, nuclear energy represents a form of firm low carbon generation characterized by high reliability and stable output. If AI expansion generates structurally higher electricity demand, and if reliability becomes a binding constraint, then the economic value of nuclear generation assets may change accordingly. The central research question of this thesis is therefore the following: Does nuclear energy represent a financially valued infrastructure for the future development of artificial intelligence? This question is addressed from a financial economics perspective. Rather than evaluating the technical feasibility of nuclear powered data centres, the analysis examines whether capital markets differentiate between electric utilities with and without nuclear exposure. If investors anticipate that AI driven electricity demand increases the strategic relevance of firm baseload generation, then nuclear exposed utilities may exhibit distinct return dynamics relative to comparable non nuclear utilities. To operationalize this question, the thesis constructs portfolios of publicly traded U.S. electric utilities classified according to their nuclear generation exposure. The empirical analysis investigates three related dimensions. First, whether nuclear exposed utilities generate statistically significant abnormal returns when evaluated under standard asset pricing models. Second, whether any performance differentials persist after controlling for common systematic risk factors. Third, whether periods characterized by heightened attention to artificial intelligence are associated with changes in relative performance between nuclear and non nuclear utilities. The empirical design is intended to test whether nuclear exposure is priced as a structural attribute in the context of AI expansion. If no differential valuation emerges, this would suggest that capital markets do not currently perceive nuclear generation as uniquely positioned within the AI driven energy transition. Conversely, evidence of abnormal performance would indicate that investors attribute economic value to nuclear assets under conditions of rising AI relevance. By framing the analysis within established asset pricing methodologies, the thesis seeks to contribute to the intersection of energy economics, infrastructure finance, and technological transition. The objective is not normative but analytical: to determine whether market based evidence supports the hypothesis that nuclear energy constitutes a strategic component of the emerging AI infrastructure.

1.4 Structure of the Thesis

The remainder of the thesis is structured as follows. Chapter 2 reviews the relevant literature at the intersection of artificial intelligence, electricity demand, and energy system economics. It first examines empirical evidence on the growth of computational requirements in frontier AI models and the associated implications for data centre electricity consumption. It then discusses the economic characteristics of different electricity generation technologies, with particular emphasis on firm low carbon resources and nuclear power. Finally, it outlines the asset pricing framework that motivates the empirical strategy. Chapter 3 presents the data and methodological approach. It describes the construction of the sample of publicly traded U.S. electric utilities and explains the classification of firms according to nuclear generation exposure. The chapter details the portfolio formation procedure, the construction of comparison groups, and the econometric specifications employed, including the Capital Asset Pricing Model and the Fama French three factor model. It also introduces an attention based extension that evaluates whether valuation patterns vary across periods of heightened artificial intelligence relevance. Chapter 4 reports the empirical results. It presents descriptive statistics, portfolio level return comparisons, and risk adjusted performance estimates. The analysis assesses whether nuclear exposed utilities exhibit abnormal returns relative to non nuclear peers and evaluates the robustness of the findings under alternative specifications. Chapter 5 discusses the interpretation of the empirical evidence. It examines whether the observed return dynamics are consistent with the hypothesis that nuclear energy represents a strategically valued infrastructure in the context of AI driven electricity demand growth. The chapter also considers alternative explanations and discusses limitations. Chapter 6 concludes by summarizing the main findings, outlining the contribution of the study to the literature on energy infrastructure and asset pricing, and suggesting avenues for future research.

1.5 Final Considerations

The expansion of artificial intelligence implies increasing reliance on electricity intensive infrastructure. As computational requirements grow, energy supply becomes a structural input into technological development. In this setting, the characteristics of electricity generation technologies may influence long term capital allocation and valuation dynamics. Nuclear energy provides firm low carbon generation with high capacity factors and continuous output. These features may become economically relevant if AI driven electricity demand increases persistently. The purpose of this thesis is to examine whether financial markets price nuclear exposure as a strategic attribute under such conditions. The analysis does not assess technological superiority or policy desirability. Instead, it evaluates whether observable return patterns are consistent with the hypothesis that nuclear generation is perceived as strategically positioned in an AI intensive economy. By grounding the discussion in asset pricing evidence, the study contributes to the broader literature on infrastructure valuation during periods of technological transition. “‘latex

2 Literature Review

2.1 The Compute–Energy Constraint of Artificial Intelligence

Large technological revolutions have historically coincided with structural transformations in energy systems. Grübler (2012) characterizes energy transitions as long-duration processes shaped by infrastructure lock-in, path dependence, and complementarities between technologies and energy carriers. Technological change is therefore embedded within energy systems rather than independent from them. The industrial revolution relied on coal, electrification required centralized grids, and the digital economy depends on scalable electricity infrastructure.

Artificial intelligence represents a new phase of digital expansion distinguished by unusually intensive computational requirements. Amodei and Hernandez (2018) document the exponential growth in compute used for frontier AI models, far exceeding historical hardware scaling trends. Kaplan et al. (2020) formalize scaling laws showing systematic performance gains with increases in parameters, data, and compute. Hoffmann et al. (2022) demonstrate that compute remains a binding input even under optimal training configurations.

Hardware developments reinforce this pattern. Jouppi et al. (2018, 2020) describe specialized accelerators and supercomputers designed for large-scale neural network workloads, emphasizing energy-performance trade-offs. Sze et al. (2017) survey efficient neural network processing architectures and highlight how improvements in efficiency coexist with rising aggregate electricity demand due to scaling.

Institutional analyses quantify macro-level implications. The International Energy Agency (2024) reports increasing electricity demand from data centers and AI workloads.

The IEA 4E (2025) review confirms upward pressure on power systems despite modeling uncertainty. Xu et al. (2025) identify grid bottlenecks associated with AI-driven data center expansion, and Nature Sustainability (2025) discusses the environmental implications of large-scale computational infrastructure. Van der Vlist, Helmond, and Ferrari (2024) characterize AI production as capital-intensive and highly centralized within cloud infrastructure. Goldman Sachs (2024) projects substantial growth in data center electricity demand over the coming decade.

The literature thus establishes that AI is structurally dependent on electricity infrastructure. Scaling compute requires sustained increases in generation capacity, linking technological progress directly to energy system constraints.

2.2 Energy Transitions and the Role of Firm Low-Carbon Capacity

Electricity system economics depend on demand profiles, generation technologies, and reliability constraints. Grübler (2012) emphasizes infrastructural inertia in energy transitions. High renewable penetration introduces variability and system balancing costs.

Hirth (2013) shows that the market value of variable renewables declines as penetration rises due to temporal correlation and price effects. Reichenberg et al. (2018) quantify system costs associated with high renewable shares, emphasizing backup and balancing requirements.

Energy storage mitigates intermittency but does not eliminate the need for firm capacity. Zakeri and Syri (2015) document heterogeneous cost effectiveness across storage technologies. The MIT Energy Initiative (2022) concludes that storage enhances flexibility but does not fully substitute dispatchable generation in deep decarbonization pathways. Mallapragada, Sepulveda, and Jenkins (2020) find that storage improves system performance but does not replace firm low-carbon resources.

System-level modeling reinforces this conclusion. Sepulveda et al. (2018) show that including firm low-carbon resources such as nuclear significantly reduces total system costs in deep decarbonization scenarios. Baik et al. (2021) demonstrate lower volatility and improved resilience in systems incorporating firm capacity.

Reliability becomes particularly relevant in the context of AI-driven electricity demand. The International Atomic Energy Agency (2025) reports growing interest in nuclear technologies to power data centers. The U.S. Department of Energy (2025) highlights high capacity factors and stable output as advantages for hyperscale computing facilities.

Cost metrics remain central to investment decisions. The International Energy Agency and OECD Nuclear Energy Agency (2020) provide cross-technology cost estimates, while Lazard (2025) presents updated levelized cost metrics. Although LCOE does not capture system value, it frames economic trade-offs across technologies.

The literature indicates that in electricity systems facing rising demand and decarbonization constraints, firm generation capacity acquires economic importance. This provides the foundation for evaluating whether equity markets differentiate between firms based on exposure to such capacity.

2.3 Nuclear Energy as Capital-Intensive Infrastructure

Nuclear generation differs structurally from both fossil fuel and variable renewable technologies. It involves high upfront capital expenditures, long construction periods, low marginal operating costs, and high capacity factors.

IEA and OECD NEA (2020) document nuclear's capital intensity and sensitivity to financing conditions. Lazard (2025) reports LCOE comparisons showing lower exposure to fuel price volatility but higher upfront costs relative to renewables. Importantly, LCOE does not incorporate system integration costs or reliability premiums.

System-level modeling emphasizes nuclear's value in decarbonized grids. Sepulveda et al. (2018) show that firm low-carbon capacity reduces total system costs under high renewable penetration. Baik et al. (2021) find improved resilience in net-zero systems incorporating firm resources.

Nuclear assets are long-duration infrastructure projects sensitive to regulatory frameworks. Werner and Jarvis (2025) revisit rate-of-return regulation and its incentive implications. Rode and Fischbeck (2019) document deviations between allowed and realized returns in regulated utilities. Michelfelder (2015) shows how regulation shapes green investment decisions. López Cabarcos et al. (2023) demonstrate that ESG characteristics contribute to explaining utility stock returns.

Operational risks include construction delays and political uncertainty. Once operational, however, nuclear plants exhibit high capacity factors, exceeding 90 percent in the United States (U.S. DOE, 2025).

Nuclear energy therefore combines infrastructure stability, regulatory exposure, and capital intensity. These features may generate distinct factor exposures relative to non-nuclear energy firms.

2.4 Asset Pricing of Utilities and Infrastructure Firms

Utility firms exhibit systematic risk exposures shaped by regulation, capital intensity, and infrastructure characteristics.

The CAPM provides a benchmark linking expected returns to market beta. Empirical evidence extends this framework through additional systematic factors, including size and value (Fama–French).

Regulated utilities display distinct risk dynamics. Rode and Fischbeck (2019) highlight deviations from traditional pricing predictions. Werner and Jarvis (2025) show that regulatory frameworks alter effective risk exposure.

Infrastructure constraints influence valuation. Kacperczyk and Licher (2025) show that electricity capacity tightness is reflected in firm value and risk premia. Xiao et al. (2024) document how energy transitions affect volatility dynamics. López Cabarcos et al. (2023) find that ESG characteristics contribute to explaining utility returns.

Nuclear utilities may therefore exhibit specific factor sensitivities linked to long asset duration, regulatory exposure, and stable baseload output. A long–short portfolio isolating nuclear minus non-nuclear exposure allows direct testing of whether return differentials persist after controlling for systematic risks.

2.5 Investor Attention, Expectations, and Google Trends

Beyond fundamental risk factors, investor attention influences asset prices. Limited attention affects how information is incorporated into markets.

Barber and Odean (2008) show that attention-driven buying behavior affects trading patterns. Da, Engelberg, and Gao (2011) introduce Google search volume as a measurable proxy for investor attention and demonstrate predictive power for short-term returns. Szczygielski et al. (2024) find that search intensity is associated with return and volatility dynamics across countries.

AI-related attention surged following the release of large-scale generative models. This shift affected expectations about infrastructure demand and electricity consumption. If markets associate nuclear capacity with AI-driven demand growth, nuclear firms may perform differently during high-attention periods.

Google Trends provides a normalized Search Volume Index (SVI). By classifying months into high and low attention regimes, it becomes possible to test whether nuclear exposure exhibits conditional return patterns.

2.6 Synthesis and Implications for the Empirical Strategy

The literature identifies three core mechanisms.

First, AI expansion increases structural electricity demand through rising computational intensity. Second, firm low-carbon resources such as nuclear provide dispatchable, stable capacity that may acquire economic relevance under demand growth and decarbonization constraints. Third, asset pricing outcomes reflect both systematic risk exposures and shifts in investor attention.

Despite these insights, limited evidence exists on whether equity markets differentiate between nuclear and non-nuclear utilities in the context of AI-related attention. In particular, it remains unclear whether nuclear exposure generates abnormal performance or whether observed differences reflect compensation for standard risk factors.

This gap motivates the empirical design. Equal-weighted nuclear and non-nuclear portfolios allow transparent comparison. Fama–French regressions assess whether performance differentials remain after controlling for systematic exposures. Google Trends–based attention regimes introduce a conditioning variable capturing shifts in narrative intensity.

The empirical analysis therefore tests whether theoretical arguments regarding compute-driven electricity demand and firm low-carbon capacity manifest in observable market pricing.

3 Data and Methodology

3.1 Firm Classification

3.1.1 Nuclear-Exposed Firms

The analysis begins with the construction of two groups of publicly traded energy companies based on structural exposure to nuclear energy. Classification is implemented manually and remains fixed throughout the sample period.

A firm is classified as nuclear-exposed if its core business is directly linked to nuclear energy activities. This includes companies that own or operate nuclear power plants, derive the majority of electricity generation from nuclear sources, or participate in the nuclear fuel cycle such as uranium production. The objective is to capture firms whose economic performance is fundamentally tied to the nuclear segment rather than firms with marginal exposure.

The nuclear-exposed sample consists of twenty publicly traded companies. The classification is static to ensure internal consistency in portfolio construction.

3.1.2 Non-Nuclear Firms

The control group consists of energy companies without exposure to nuclear activities. A firm is classified as non-nuclear if it does not own or operate nuclear power plants, does not generate electricity from nuclear sources, and does not participate in uranium extraction or related activities.

These firms operate in conventional generation, renewables, fossil fuels, or energy infrastructure services. The non-nuclear sample also contains twenty publicly traded companies, matching the nuclear group to maintain symmetry in portfolio construction and facilitate direct comparison.

3.2 Equity Data and Return Construction

3.2.1 Price Data

Monthly adjusted closing prices are obtained using the `yfinance` library for all forty firms over the period January 2016 to January 2026. Adjusted prices incorporate dividends and corporate actions, ensuring total return measurement.

Data are retrieved at monthly frequency and aligned to month-end timestamps to ensure consistency with the Fama–French factor data used in the regression analysis. Only firms with valid price data over the specified horizon are retained. No external filtering or imputation procedures are applied.

3.2.2 Return Computation

Monthly returns are computed using simple percentage changes:

$$R_{i,t} = \frac{P_{i,t} - P_{i,t-1}}{P_{i,t-1}}, \tag{3.1}$$

where $P_{i,t}$ denotes the adjusted closing price of firm i at month-end t.

Missing observations are not interpolated. The first observation for each firm is removed due to lack of a lagged price. All returns are aligned to month-end timestamps to ensure consistency across datasets.

3.3 Portfolio Construction

3.3.1 Equal-Weighted Nuclear Portfolio

The nuclear portfolio is constructed using equal weighting. For each month t:

$$R_t^{Nuclear} = \frac{1}{N_t} \sum_{i=1}^{N_t} R_{i,t}, \tag{3.2}$$

where N_t denotes the number of nuclear firms with available returns in month t.

Each firm contributes equally. Firms with missing returns in a given month are excluded from that month's average.

3.3.2 Equal-Weighted Non-Nuclear Portfolio

The non-nuclear portfolio is constructed analogously:

$$R_t^{NonNuclear} = \frac{1}{M_t} \sum_{j=1}^{M_t} R_{j,t}, \tag{3.3}$$

where M_t is the number of non-nuclear firms with available returns in month t.

3.3.3 Long–Short Factor Construction

To isolate relative performance, a self-financing long–short factor is defined:

$$NLS_t = R_t^{Nuclear} - R_t^{NonNuclear}. \tag{3.4}$$

The Nuclear Long–Short factor (NLS) represents a strategy that is long nuclear and short non-nuclear firms with equal notional exposure.

3.4 Fama–French Three-Factor Regression Framework

3.4.1 Factor Data

Monthly Fama–French three-factor data (Mkt-RF, SMB, HML, and RF) are imported from a local CSV file. Dates are converted to month-end timestamps and aligned with portfolio returns. Factor values expressed in percentage format are rescaled to decimals when necessary.

The merged dataset includes only months with complete observations.

3.4.2 Regression Specifications

The Nuclear Long–Short factor is estimated using:

$$NLS_t = \alpha + \beta_M (Mkt\text{-}RF)_t + \beta_S SMB_t + \beta_H HML_t + \varepsilon_t. \tag{3.5}$$

Since the NLS factor is self-financing, the dependent variable is not adjusted by subtracting the risk-free rate.

For the long-only portfolios:

$$(R_t^p - RF_t) = \alpha^p + \beta_M^p (Mkt\text{-}RF)_t + \beta_S^p SMB_t + \beta_H^p HML_t + \varepsilon_t^p, \tag{3.6}$$

where $p \in \{Nuclear, NonNuclear\}$.

3.4.3 Estimation Procedure

All regressions are estimated using Ordinary Least Squares. For time-series regressions on portfolio returns, conventional OLS standard errors are reported. Coefficient estimates, standard errors, R^2, adjusted R^2, and the number of observations are presented.

3.5 Performance Metrics

In addition to regression analysis, portfolio performance is evaluated using descriptive statistics computed from monthly returns.

The annualized mean return is obtained by multiplying the monthly mean by 12. Annualized volatility is computed as the monthly standard deviation multiplied by $\sqrt{12}$.

Cumulative return indices are defined as:

$$Index_t = \prod_{s=1}^{t} (1 + R_s), \tag{3.7}$$

with initial value equal to 1.

Drawdowns are computed as:

$$DD_t = \frac{Index_t}{\max_{s \le t} Index_s} - 1. \tag{3.8}$$

The maximum drawdown is defined as the minimum of DD_t over the sample period.

3.6 Google Trends and Attention-Based Analysis

3.6.1 Google Trends Data

Monthly Google Trends Search Volume Index (SVI_t) for artificial intelligence is collected and aligned to month-end timestamps.

3.6.2 Attention Regimes

Attention regimes are defined using sample quantiles:

$$Regime_t = \begin{cases} Low, & SVI_t \leq q_{0.30}, \\ High, & SVI_t \geq q_{0.70}, \\ Mid, & \text{otherwise.} \end{cases} \tag{3.9}$$

High and Low months are used for conditional mean comparisons.

3.6.3 High–Low Return Differences

For each return series R_t:

$$\Delta\mu(R) = \mathbb{E}[R_t \mid High] - \mathbb{E}[R_t \mid Low]. \tag{3.10}$$

Annualized differences are obtained by multiplying monthly averages by 12.

3.6.4 Pooled Regression

Firm-level excess returns are modeled as:

$$ExcessR_{i,t} = \alpha + \beta_1 Nuclear_i + \beta_2 SVI_t + \beta_3(Nuclear_i \times SVI_t) + \varepsilon_{i,t}. \tag{3.11}$$

Standard errors are clustered by month to account for cross-sectional dependence.

3.7 Methodological Considerations

Firms are classified statically based on core business activities. This ensures consistency but does not account for potential changes in generation mix over time.

Equal weighting prevents large firms from dominating portfolio returns but may amplify the influence of smaller firms relative to capitalization-weighted benchmarks.

Monthly frequency reduces microstructure noise and aligns with factor data availability, though it may smooth short-lived market reactions.

Google Trends is used as a proxy for attention rather than direct capital flows. Regime thresholds are sample-dependent and reflect relative attention intensity within the observed period.

These design choices aim to balance transparency, comparability, and empirical tractability.

4 Empirical Results

4.1 Descriptive Performance of Portfolios

4.1.1 Return and Volatility

Table 4.1 reports the main descriptive statistics for the equal-weighted Nuclear portfolio, the equal-weighted Non-Nuclear portfolio, and the Nuclear Long–Short factor (NLS) over the sample period from January 2016 to January 2026, comprising 119 monthly observations.

The Non-Nuclear portfolio exhibits the highest average performance, with a monthly mean return of 0.018, corresponding to an annualized mean of 0.214. The Nuclear portfolio records a monthly mean return of 0.014, equivalent to an annualized return of 0.166. The NLS factor produces a negative average monthly return of -0.004, corresponding to an annualized return of -0.049.

In terms of volatility, the Non-Nuclear portfolio displays the highest monthly standard deviation at 0.062, followed by the Nuclear portfolio at 0.048. The NLS factor exhibits a monthly volatility of 0.044. Maximum drawdowns indicate that the Nuclear portfolio experienced a peak-to-trough decline of -0.228, while the Non-Nuclear portfolio reached -0.284. The NLS factor shows a substantially larger maximum drawdown of -0.671, reflecting the risk profile of the spread strategy.

Unconditionally, nuclear exposure does not appear to deliver a persistent return premium over the full sample. If nuclear energy were already priced as indispensable infrastructure for AI driven electricity demand, markets would be expected to embed forward looking premia in nuclear exposed firms, which would translate into sustained

Table 4.1: Descriptive Performance Statistics

	Mean (Monthly)	Volatility (Monthly)	Mean (Annualized)	Max Drawdown
NuclearEW	0.014	0.048	0.166	-0.228
NonNuclearEW	0.018	0.062	0.214	-0.284
NLS	-0.004	0.044	-0.049	-0.671

relative outperformance or, at minimum, a non negative average long short spread. Instead, the negative mean of NLS suggests relative underperformance of the nuclear group.

At the same time, NuclearEW exhibits lower volatility and a smaller maximum drawdown than NonNuclearEW. This pattern is consistent with nuclear utilities behaving as more defensive regulated assets with comparatively stable cash flows, rather than as high growth infrastructure beneficiaries. The combination of weaker average performance and lower downside risk motivates the risk adjusted analysis and the attention based tests that follow, since any AI related valuation effect may be state dependent rather than persistent across the entire sample.

An alternative perspective on Table 4.1 emerges from comparing the relative magnitude of return and risk differentials. While NonNuclearEW generates a monthly return that exceeds NuclearEW by 0.004, the associated increase in dispersion amounts to 0.014. The return differential is therefore small relative to the additional variability borne by investors. From a portfolio construction standpoint, nuclear exposure appears weaker in absolute terms but not dramatically inferior once risk is taken into account. This suggests that nuclear firms were positioned in a lower beta segment of the sector rather than being systematically penalized by markets.

4.1.2 Cumulative Returns

Figure 4.1 reports the cumulative performance (base 1) of NuclearEW and NonNuclearEW.

The cumulative return paths indicate that both portfolios generate substantial long-term growth over the sample period. However, the Non-Nuclear portfolio consistently outperforms the Nuclear portfolio in compounded terms, leading to a higher terminal wealth level. The divergence between the two series becomes particularly visible during periods of strong aggregate equity performance, when the higher market beta of the

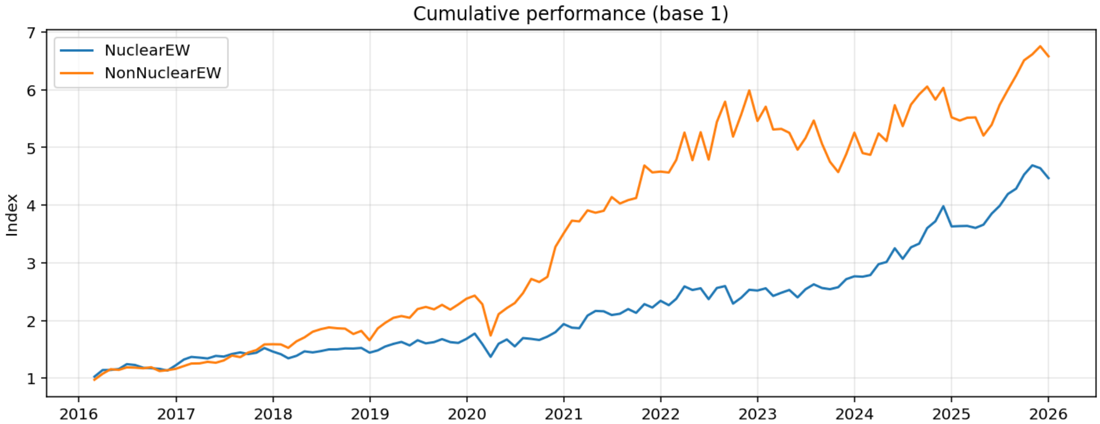

Figure 4.1: Cumulative performance (base 1) of NuclearEW and NonNuclearEW, January 2016 to January 2026.

Non-Nuclear portfolio translates into larger cumulative gains.

This persistent compounding gap is difficult to reconcile with the view that nuclear exposure was systematically valued as a long run AI infrastructure beneficiary over the sample. If such a narrative had been continuously embedded in prices, nuclear exposed firms would be expected to match or exceed the performance of their counterparts. Instead, relative pricing favored the non nuclear set precisely when market conditions rewarded cyclical risk taking, which aligns more closely with conventional factor exposure than with a durable AI linked premium.

The timing of the divergence also indicates that relative underperformance is not uniform across states. The gap widens when risk appetite is elevated and narrows during more defensive market phases. This asymmetry supports the interpretation that nuclear firms were not treated as technological growth accelerators, but rather as lower beta utilities whose valuations respond less aggressively to broad market upswings.

4.1.3 Drawdown Analysis

Figure 4.2 reports the drawdowns of NuclearEW and NonNuclearEW.

The drawdown profiles reveal that although the Non-Nuclear portfolio achieves higher cumulative performance, it also experiences deeper peak-to-trough losses. The Nuclear portfolio exhibits comparatively milder drawdowns, suggesting a more defensive profile during adverse market conditions. This pattern is consistent with the lower market beta

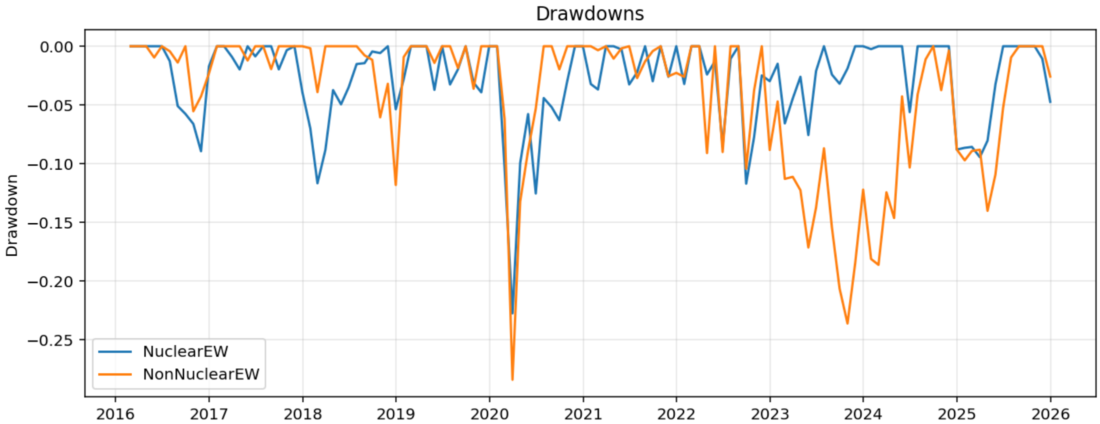

Figure 4.2: Drawdowns of NuclearEW and NonNuclearEW, January 2016 to January 2026.

and volatility documented in Table 4.1.

The drawdown evidence reinforces the view that nuclear exposure is associated with a more defensive risk profile. This is compatible with nuclear firms being priced primarily through regulated utility fundamentals and stability considerations. If the AI electricity narrative affects nuclear valuation, this section suggests it does not dominate pricing continuously, which makes a conditional channel more plausible.

4.1.4 NLS Cumulative Performance and Drawdowns

Figure 4.3 reports the cumulative performance (base 1) of NuclearEW, NonNuclearEW, and the Nuclear Long–Short factor (NLS).

The NLS factor displays substantially weaker long-run compounding relative to both long-only portfolios. Its cumulative trajectory fluctuates around a declining trend, reflecting the negative average return reported earlier. This indicates that, over the sample period, nuclear exposure underperformed non-nuclear exposure on a relative basis.

Figure 4.4 reports drawdowns for all three series.

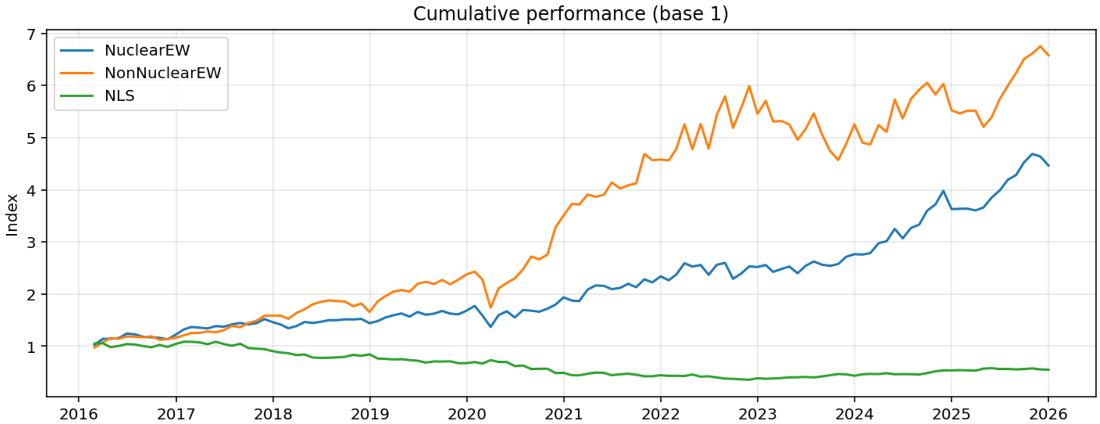

Figure 4.3: Cumulative performance (base 1) of NuclearEW, NonNuclearEW, and NLS, January 2016 to January 2026.

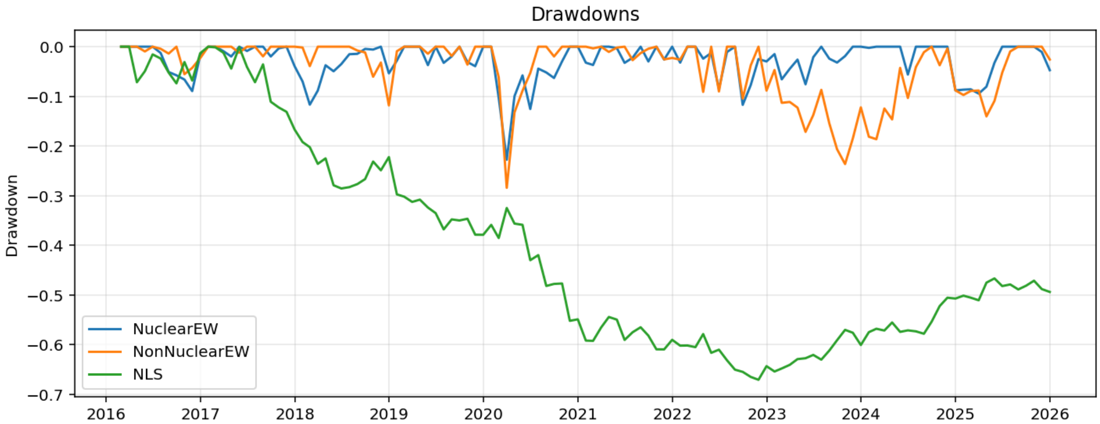

Figure 4.4: Drawdowns of NuclearEW, NonNuclearEW, and NLS, January 2016 to January 2026.

The NLS factor exhibits markedly larger and more persistent drawdowns compared to the long-only portfolios. This reflects the cumulative effect of relative underperformance and highlights the risk concentration inherent in a spread strategy. While both underlying portfolios recover from downturns over time, the NLS factor experiences prolonged negative phases driven by sustained relative gaps between nuclear and non-nuclear returns.

Taken together, Section 4.1 provides no evidence of an unconditional nuclear premium. The descriptive patterns are consistent with nuclear exposure behaving as a defensive utility characteristic rather than as a persistent AI growth premium. This does not exclude episodic repricing. A natural hypothesis is that valuation effects emerge when nuclear energy becomes salient in public and policy discourse, which is investigated in Section 4.3.

4.2 Fama–French Three-Factor Results

This section reports the Fama–French three-factor regressions estimated in Chapter 3. The NuclearEW and NonNuclearEW portfolios are analyzed in excess-return form, while the Nuclear Long–Short factor (NLS) is analyzed in raw-return form because it is self-financing. The sample spans January 2016 to January 2026 with 119 monthly observations.

4.2.1 Regression estimates

Table 4.2 reports coefficient estimates, standard errors, and goodness-of-fit statistics. For NuclearEW and NonNuclearEW, the dependent variable is $R_t - RF_t$. For NLS, the dependent variable is NLS_t.

Table 4.2: Fama–French Three-Factor Regression Results

	NLS	NuclearEW (Excess)	NonNuclearEW (Excess)
α	-0.002	0.004	0.006^{*}
	(0.004)	(0.003)	(0.004)
$Mkt-RF$	-0.242^{***}	0.715^{***}	0.957^{***}
	(0.080)	(0.079)	(0.080)
SMB	-0.612^{***}	-0.213^{*}	0.399^{***}
	(0.131)	(0.129)	(0.131)
HML	-0.131	0.186^{**}	0.317^{***}
	(0.089)	(0.088)	(0.089)
R^2	0.298	0.435	0.651
Adj. R^2	0.280	0.420	0.642
N	119	119	119

Notes. Standard errors are in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively. NuclearEW and NonNuclearEW are modeled as monthly excess returns $R_t - RF_t$. NLS is modeled in raw returns because it is a self-financing spread.

4.2.2 Interpretation of factor exposures

The market factor loading is positive and strongly significant for both long-only portfolios, with a larger exposure for the NonNuclearEW portfolio. In contrast, the NLS factor has a negative and statistically significant loading on $Mkt-RF$, consistent with the spread behaving as a negative market exposure on average.

The size factor loadings differ sharply across portfolios. The NuclearEW portfolio has a negative SMB loading, while the NonNuclearEW portfolio has a positive SMB loading. The NLS factor inherits a strongly negative and statistically significant SMB exposure, indicating that the return spread is associated with the size dimension captured by the Fama–French factors.

For the value factor, the NuclearEW and NonNuclearEW portfolios both load positively on HML, with stronger statistical support for the NonNuclearEW portfolio. The NLS loading on HML is negative but not statistically significant at conventional levels.

These results indicate that relative performance differences are largely attributable to systematic risk exposures. In particular, the absence of a statistically significant positive alpha for NuclearEW implies that nuclear exposure does not generate persistent abnormal returns beyond conventional risk compensation. If nuclear energy were already priced as foundational AI infrastructure, forward looking premia would be expected to translate into a robust positive alpha. The evidence does not support this. Instead, nuclear exposure appears to be valued through conventional factor channels, while any AI related repricing is more likely to operate conditionally, which motivates the attention based analysis that follows.

The negative and significant market loading of NLS implies that the relative return of nuclear versus non-nuclear firms is systematically related to aggregate market conditions. When the market performs strongly, the spread tends to move against nuclear exposure. This indicates that relative underperformance is not purely idiosyncratic but embedded in broader risk dynamics. Any interpretation of nuclear valuation therefore needs to distinguish between structural repricing and compensation for systematic exposure differences.

4.3 Attention Regimes and Google Trends Evidence

This section evaluates whether public attention to nuclear energy, proxied by Google Search Volume Index (SVI), affects the relative performance of nuclear and non-nuclear utilities.

Monthly observations are classified into regimes using the 30th and 70th percentiles of the SVI distribution. Months below the 30th percentile are defined as Low attention, while months above the 70th percentile are defined as High attention. The High and Low regimes contain 37 and 55 observations, respectively.

4.3.1 High versus Low Attention Regimes

Table 4.3 reports average monthly returns conditional on attention regimes.

Table 4.3: Mean Monthly Returns by Attention Regime

	High SVI	Low SVI	Difference (High–Low)
NuclearEW	0.016	0.010	0.006
NonNuclearEW	0.004	0.018	-0.014
NLS	0.012	-0.008	0.020

The Nuclear portfolio performs better during periods of elevated public attention, while the Non-Nuclear portfolio performs substantially better during low-attention periods. The Nuclear Long–Short factor delivers positive returns during High attention regimes and negative returns during Low attention regimes, implying a return differential of approximately 2 percentage points per month.

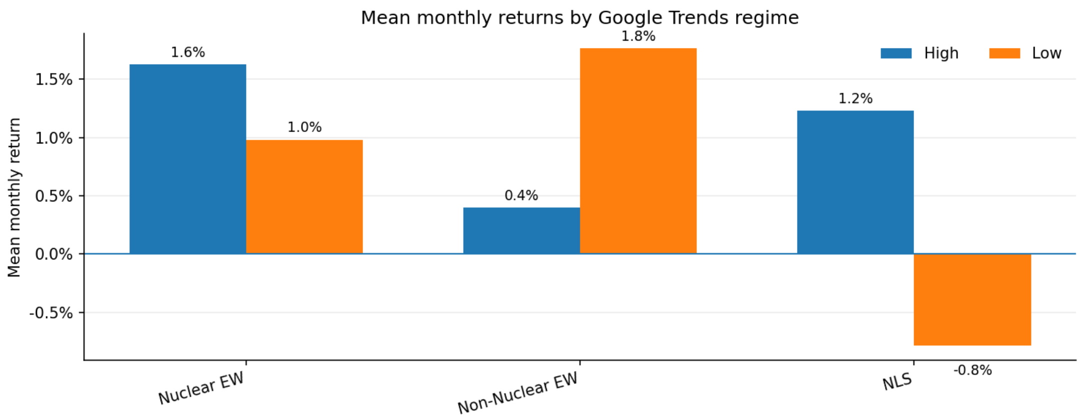

Figure 4.5: Mean monthly returns across High and Low attention regimes.

The bar plot visually highlights the regime-dependent reversal between the Nuclear and Non-Nuclear portfolios.

To assess economic magnitude, monthly returns are annualized.

Table 4.4: Annualized Mean Returns by Attention Regime

	High SVI	Low SVI	Difference (High–Low)
NuclearEW	0.195	0.118	0.077
NonNuclearEW	0.048	0.212	-0.164
NLS	0.147	-0.094	0.242

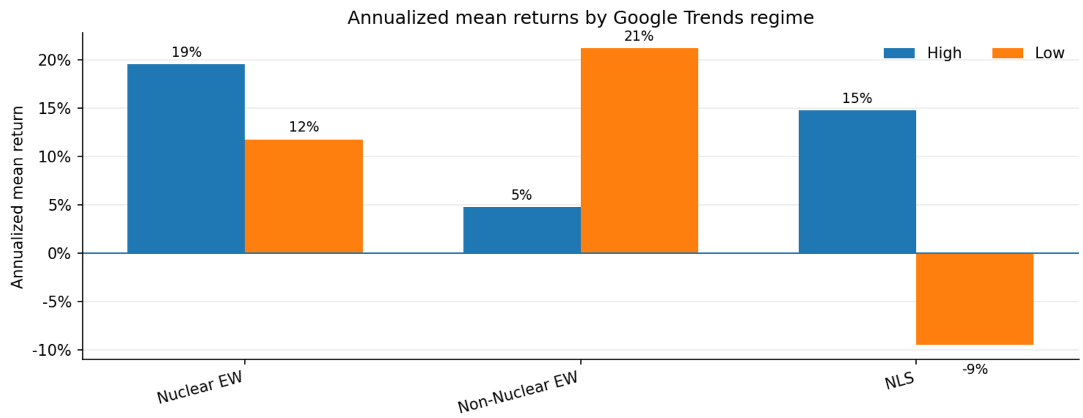

Figure 4.6: Annualized mean returns across High and Low attention regimes.

The economic magnitude is substantial. The NLS factor delivers an annualized return spread of approximately 24 percentage points between High and Low attention regimes. This suggests that public attention materially alters relative performance within the utility sector.

These regime results provide the clearest evidence in the thesis that nuclear exposure is priced conditionally. Nuclear outperformance does not appear as a continuous phenomenon, but it emerges when nuclear energy becomes salient in public discourse. This is consistent with attention operating as a transmission channel between macro narratives and equity pricing. In the context of AI driven electricity demand, the results suggest that nuclear exposure is repriced when the energy narrative intensifies, without implying a permanent valuation premium across all states.

The magnitude of the attention effect is large relative to unconditional averages. The annualized High–Low spread for NLS of 24 percentage points substantially exceeds the unconditional annualized mean of -0.049 reported in Table 4.1. This indicates that regime shifts are not marginal adjustments around a stable mean, but represent economically meaningful changes in relative valuation across states.

A structural shift would imply sustained abnormal performance and persistent repricing. The evidence here is more consistent with cycclical repricing, since relative performance shifts with attention and reverses when attention is low. The broader structural versus cyclical interpretation, and its implications for whether nuclear energy becomes foundational AI infrastructure, is developed further in the conclusion.

4.3.2 Pooled Regression with Attention Interaction

To formally test whether attention modifies the relative performance of nuclear firms, the following pooled regression is estimated:

$$ExcessR_{i,t} = \alpha + \beta_1 Nuclear_i + \beta_2 SVI_t + \beta_3 (Nuclear_i \times SVI_t) + \varepsilon_{i,t} \tag{4.1}$$

where $ExcessR_{i,t}$ denotes firm-level excess returns, $Nuclear_i$ is a dummy variable equal to one for nuclear-exposed firms, and SVI_t measures public attention in month t. Standard errors are clustered at the monthly level.

Table 4.5: Pooled OLS with Attention Interaction

Variable	Coefficient	Std. Error
Intercept	0.020^{***}	(0.007)
Nuclear	-0.009^{*}	(0.005)
SVI	-0.0002	(0.0002)
Nuclear × SVI	0.0003^{**}	(0.0001)
R^2	0.002	
Observations	4,751	

Standard errors are clustered by month. ***, **, * denote significance at the 1%, 5%, and 10% levels.

The negative coefficient on the Nuclear dummy suggests that, on average, nuclear firms underperform non-nuclear firms. However, the interaction term between Nuclear exposure and SVI is positive and statistically significant. This indicates that higher public attention mitigates or reverses the relative underperformance of nuclear firms.

Economically, the positive interaction coefficient implies that the return differential between nuclear and non-nuclear firms increases as attention rises. This regression evidence confirms the regime-based findings presented in Section 4.3.1 and provides formal statistical support for the attention effect.

The interaction term also has a cross-sectional interpretation. A positive Nuclear $\times$ SVI coefficient implies that the relative pricing of nuclear firms becomes more favorable precisely when attention rises, even after controlling for average nuclear underperformance. This suggests that the market response operates at the firm level rather than purely at the portfolio aggregation level. The consistency between regime averages and firm-level regressions strengthens the interpretation that attention affects valuation directly, rather than through mechanical portfolio composition effects.

Taken together with the regime results, the pooled regression indicates that nuclear valuation is attention sensitive. Nuclear exposure does not command a persistent premium, but it becomes relatively more valuable when attention rises. This supports a nuanced answer to the research question: nuclear exposure is not priced as a permanently embedded AI infrastructure premium during the sample period, but it exhibits economically meaningful state dependence consistent with episodic repricing when AI energy discourse and related nuclear narratives become salient. Limits of the attention proxy, the binary exposure classification, and the indirect nature of linking AI electricity demand to utility pricing are addressed in the conclusion, together with the broader interpretation of whether these dynamics reflect early stages of a structural shift or cyclical repricing.

5 Conclusion

5.1 Summary of Findings and Answer to the Research Question

This thesis studied whether companies involved in nuclear power generation gain a valuation advantage as AI-driven electricity demand expands. The central question was whether capital markets price nuclear exposure as foundational infrastructure for the expansion of large-scale computational activity.

The empirical evidence does not indicate the existence of a persistent nuclear premium over the sample period from January 2016 to January 2026. Unconditional portfolio comparisons show that nuclear-exposed utilities underperform their non-nuclear counterparts in compounded terms, and the long–short spread yields a negative average return. Risk-adjusted analysis further demonstrates that this relative performance difference is largely explained by systematic factor exposures. In particular, nuclear exposure does not generate statistically significant positive alpha once conventional risk factors are controlled for.

At the same time, the attention-based analysis reveals pronounced state dependence. During periods of elevated public attention to nuclear energy, nuclear-exposed firms outperform on a relative basis, and the long–short factor delivers economically meaningful positive returns. When attention is low, the performance differential reverses. Firm-level regressions confirm that the interaction between nuclear exposure and attention is positive and statistically significant, indicating that the effect operates beyond simple portfolio aggregation.

Taken together, the findings suggest that nuclear exposure is not priced as a persistently embedded AI infrastructure premium. Instead, its relative valuation fluctuates systematically with changes in narrative intensity. The evidence therefore supports a conditional, rather than structural, interpretation of nuclear repricing within the observed sample.

5.2 Interpretation and Broader Implications

The absence of a persistent abnormal return implies that markets have not yet treated nuclear energy as a continuously rewarded growth input for AI development. If investors had fully internalized a long-run expectation of structurally rising electricity demand tied to AI expansion, nuclear-exposed firms would be expected to exhibit sustained relative strength or positive risk-adjusted performance. This pattern does not emerge in the data.

However, the regime results indicate that valuation responds meaningfully when nuclear energy becomes salient in public discourse. The magnitude of the attention-driven return differential is large relative to unconditional averages, suggesting that narrative intensity can temporarily alter capital allocation within the utility sector. This behavior is consistent with markets reacting to shifts in perceived strategic relevance rather than continuously pricing in long-term structural transformation.

An additional implication concerns timing. The relative underperformance of nuclear exposure is concentrated during periods of strong aggregate market performance, while its defensive characteristics become more visible during less risk-seeking environments. This asymmetry suggests that nuclear assets are not treated as high-growth technological complements to AI, but rather as regulated infrastructure with limited participation in broad equity rallies. The attention effect therefore appears as episodic repricing layered on top of a fundamentally defensive risk profile.

The combined evidence points toward cyclical repricing rather than completed structural revaluation. Nuclear exposure becomes relatively more attractive when attention to energy security, decarbonization, or AI-related electricity demand intensifies, but this adjustment does not persist across states. Whether such episodes represent early signals of a gradual structural shift remains unresolved. Within the sample period, however, the pricing dynamics are better characterized as state-dependent adjustments rather than permanent incorporation of an AI infrastructure premium.

5.3 Limitations and Future Research

Several limitations qualify the interpretation of the results. The attention proxy captures public search intensity rather than direct institutional capital flows, and therefore measures narrative salience rather than explicit investment commitments. The binary classification of firms into nuclear and non-nuclear categories abstracts from heterogeneity in generation mix intensity and strategic positioning. In addition, the analysis focuses on electricity producers rather than on direct contractual relationships between AI firms and specific energy suppliers.

Future research may refine the identification strategy by incorporating measures of generation capacity shares, capital expenditure in nuclear infrastructure, or contractual exposure to data center operators. Extending the analysis to longer horizons or incorporating forward-looking valuation metrics may help determine whether episodic repricing evolves into persistent structural transformation. Linking firm-level energy procurement data from AI-intensive sectors to utility revenue streams could also provide more direct evidence of infrastructure dependence.

Within the scope of the present study, the findings indicate that nuclear energy exposure is not yet priced as a permanent AI growth characteristic, but exhibits economically meaningful state dependence aligned with shifts in attention. The distinction between conditional repricing and structural transformation remains central for understanding how energy infrastructure is incorporated into equity valuation in an era of accelerating technological demand.

References

Amodei, D., & Hernandez, D. (2018). AI and compute. OpenAI. `https://openai.com/index/ai-and-compute/`

Baik, E., Jenkins, J. D., Sepulveda, N. A., & Lester, R. K. (2021). What's different about different net-zero electricity systems? *Energy & Climate Change, 2*, 100046.

Barber, B. M., & Odean, T. (2008). All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *Review of Financial Studies, 21*(2), 785–818.

Da, Z., Engelberg, J., & Gao, P. (2011). In search of attention. *Journal of Finance, 66*(5), 1461–1499.

Goldman Sachs. (2024). *AI to drive 165% increase in data center power demand by 2030.*

Grübler, A. (2012). Energy transitions research: Insights and cautionary tales. *Energy Policy, 50*, 8–16.

Hirth, L. (2013). The market value of variable renewables: The effect of solar and wind power variability. *Energy Economics, 38*, 218–236.

Hoffmann, J., Borgeaud, S., Mensch, A., et al. (2022). Training compute-optimal large language models. *Advances in Neural Information Processing Systems.* `https://arxiv.org/abs/2203.15556`

International Atomic Energy Agency. (2025). *Data centres, artificial intelligence and cryptocurrencies eye advanced nuclear to meet growing power demand.*

International Energy Agency. (2024). *Energy and AI: Energy demand from AI.* `https://www.iea.org/reports/energy-and-ai`

International Energy Agency 4E. (2025). *Data centre energy use: Critical review of models and results.*

International Energy Agency & OECD Nuclear Energy Agency. (2020). *Projected costs of generating electricity 2020.* OECD Publishing.

Jouppi, N. P., Young, C., Patil, N., & Patterson, D. A. (2018). Motivation for and evaluation of the first tensor processing unit. *IEEE Micro, 38*(3), 10–19.

Jouppi, N. P., Yoon, D. H., Kurian, G., et al. (2020). A domain-specific supercomputer for training deep neural networks. *Communications of the ACM, 63*(7), 67–78.

Kacperczyk, M., & Licher, D. (2025). *Infrastructure capacity, risk, and firm value: Evidence from U.S. electricity tightness* (HKU Jockey Club Enterprise Sustainability Global Research Institute Paper Series, 2025/136).

Kaplan, J., McCandlish, S., Henighan, T., et al. (2020). Scaling laws for neural language models. *arXiv preprint* arXiv:2001.08361. `https://arxiv.org/abs/2001.08361`

Lazard. (2025). *Levelized cost of energy+ (LCOE+) analysis* (Version 19.0). Lazard.

López-Cabarcos, M. Á., Santos-Rodríguez, H., Quiñoá-Piñeiro, L., & Piñeiro-Chousa, J. (2023). How to explain stock returns of utility companies from an environmental, social and corporate governance perspective. *Corporate Social Responsibility and Environmental Management.*

MacKinlay, A. C. (1997). Event studies in economics and finance. *Journal of Economic Literature, 35*(1), 13–39.

Mallapragada, D. S., Sepulveda, N. A., & Jenkins, J. D. (2020). Long-run system value of battery energy storage in future grids with increasing wind and solar generation. *Applied Energy, 275*, 115390.

Massachusetts Institute of Technology. (2022). *The future of energy storage: An interdisciplinary MIT study.* MIT Energy Initiative.

Michelfelder, R. A. (2015). Electric utility regulation and investment in green energy resources. *Journal of Sustainable Finance & Investment, 5*(1–2), 48–64.

Nature Sustainability. (2025). Environmental impacts of AI data centres.

Reichenberg, L., et al. (2018). The system cost of variable renewables: Baseload trade-offs in European power systems. *Renewable Energy.*

Rode, D. C., & Fischbeck, P. S. (2019). Regulated equity returns: A puzzle. *Energy Policy, 134*, 110948.

Sepulveda, N. A., Jenkins, J. D., de Sisternes, F. J., & Lester, R. K. (2018). The role of firm low-carbon electricity resources in deep decarbonization of power generation. *Joule, 2*(11), 2403–2420.

Szczygielski, J. J., Charteris, A., Bwanya, P. R., & Brzeszczyński, J. (2024). Google search trends and stock markets: Sentiment, attention or uncertainty? *International Review of Financial Analysis, 91*, 102549.

Sze, V., Chen, Y.-H., Yang, T.-J., & Emer, J. S. (2017). Efficient processing of deep neural networks: A tutorial and survey. *Proceedings of the IEEE, 105*(12), 2295–2329.

U.S. Department of Energy. (2025). *Advantages and challenges of nuclear-powered data centers.*

Van der Vlist, F. N., Helmond, A., & Ferrari, F. (2024). Big AI: Cloud infrastructure dependence and the industrialisation of artificial intelligence. *Big Data & Society, 11*(1).

Werner, K. D., & Jarvis, S. (2025). *Rate of return regulation revisited* (Energy Institute at Haas Working Paper 329R).

Xiao, X., Li, A., Kchouri, B., & Shan, S. (2024). Tracing the dynamic impact of energy transitions on equity market volatility in an era of financial turbulence. *Energy Economics, 133*, 107443.

Xu, Y., et al. (2025). *Electricity demand and grid impacts of AI data centers.* arXiv.

Zakeri, B., & Syri, S. (2015). Electrical energy storage systems: A comparative life cycle cost analysis. *Renewable and Sustainable Energy Reviews, 42*, 569–596.

A Classification of Nuclear and Non-Nuclear Firms

This appendix reports the classification of firms used in the empirical analysis. Companies are divided into nuclear-exposed and non-nuclear firms.

A.1 Nuclear-Exposed Firms

Table A.1: Nuclear-Exposed Tickers and Company Names

Ticker	**Company Name**
VST	Vistra Corp.
ETR	Entergy Corporation
DUK	Duke Energy Corporation
D	Dominion Energy, Inc.
SO	The Southern Company
PEG	Public Service Enterprise Group Inc.
NEE	NextEra Energy, Inc.
PCG	PG&E Corporation (Pacific Gas and Electric)
AEP	American Electric Power Company, Inc.
DTE	DTE Energy Company
AEE	Ameren Corporation
XEL	Xcel Energy Inc.
EVRG	Evergy, Inc.
PNW	Pinnacle West Capital Corporation
CCJ	Cameco Corporation
BWXT	BWX Technologies, Inc.
EIX	Edison International
TXNM	TXNM Energy, Inc. (formerly PNM Resources, Inc.)
UEC	Uranium Energy Corp.
NXE	NexGen Energy Ltd.

A.2 Non-Nuclear Firms

Table A.2: Non-Nuclear Tickers and Company Names

Ticker	Company Name
NRG	NRG Energy, Inc.
FSLR	First Solar, Inc.
ENPH	Enphase Energy, Inc.
ORA	Ormat Technologies, Inc.
CWEN	Clearway Energy, Inc.
LNT	Alliant Energy Corporation
AES	The AES Corporation
ED	Consolidated Edison, Inc.
SRE	Sempra
NI	NiSource Inc.
CMS	CMS Energy Corporation
ATO	Atmos Energy Corporation
OKE	ONEOK, Inc.
WMB	The Williams Companies, Inc.
LNG	Cheniere Energy, Inc.
DVN	Devon Energy Corporation
COP	ConocoPhillips
EOG	EOG Resources, Inc.
RUN	Sunrun Inc.
PLUG	Plug Power Inc.