De-Risking Nuclear Energy Investment is NEAR with Financial Innovation

Saya Lee
Pennsylvania State University
email: sayalee@psu.edu

Jilang Miao
Pennsylvania State University
email: jlmiao@psu.edu

Nonna Sorokina*
Pennsylvania State University
Phone: 570-963-2662
email: sorokina@psu.edu

Darshana Sunoj
Pennsylvania State University
email: darshanasunoj@gmail.com

This Draft: February 27, 2026

Corresponding author:

Nonna Sorokina, email sorokina@psu.edu, ph. 570-963-2662

The research project is generously funded by the Nuclear Energy University Program (NEUP), U.S. Department of Energy grant number is DE-NE0009381, Penn State Institute for Computational and Data Sciences (RRID: SCR_025154), and Penn State University Center for Applications of Artificial Intelligence and Machine Learning to Industry Core Facility (RRID: SCR_022867). Guangxian Song was also supported via the Institute's Rising Scholar initiative. We are incredibly thankful to Gretta Kellogg for her continued support. Research assistantships of Janice Tran, Tanuja Voruganti, Santhosh Boominathan, and Guangxian Song was phenomenal and vital for the success of this project.

De-Risking Nuclear Energy Investment is NEAR with Financial Innovation

Abstract

Nuclear energy stands as a pivotal pillar for achieving energy independence, economic leadership, and long-term sustainability, yet private investment has lagged far behind its strategic importance. We introduce the Nuclear Energy Acceptance Rating (NEAR), a novel sentiment-based measure that translates societal attitudes into a finance-ready signal for investors. Grounded in the Sociotechnical Readiness Level framework (Verma & Allen, 2024) and extended through sentiment-informed advances (Kim et al., 2025[a], 2025[b]), NEAR captures evolving public acceptance and integrates it into a rolling mean–variance portfolio optimizer with option-based hedging. This framework combines the diversification logic of megafunds with the targeted risk-transfer of FDA hedges, creating an instrument capable of transforming nuclear finance. Preliminary results are promising; baseline result of the proposed hedged portfolio generated 12.5 percent improvement in risk-adjusted return over equally weighted equity-only portfolio even when our period included periods of financial and industry specific crises, including financial crisis of 2007-2009 and post-Fukushima tensions in the nuclear energy industry. By embedding societal readiness into financial design, NEAR unlocks the potential to flood nuclear energy with private capital at unprecedented scale.

JEL codes: G11, G17, Q56

Keywords: Sustainability, Portfolio, Derivatives, Risk Management, Energy, Investment, Sentiment, NLP, Nuclear, Decarbonization.

De-Risking Nuclear Energy Investment is NEAR with Financial Innovation

"The Stone Age Didn't End Because We Ran Out of Stones."

Thomas Friedman

1 Introduction.

Financing nuclear energy remains one of the most persistent challenges in the global transition to low-carbon power. Despite its capacity to provide stable, carbon-free baseload electricity, nuclear investment continues to be constrained by high capital intensity, long development horizons, regulatory complexity, and societal uncertainty. Weibezahn and Steigerwald (2024) document that nearly all nuclear projects in Europe rely on direct or indirect government support, reflecting limited private capital participation. El Ghoul et al. (2011) show that firms operating in controversial industries face systematically higher costs of equity, indicating that reputational and regulatory risks are embedded in market pricing. Even emerging technologies such as Small Modular Reactors face commercialization, licensing, security, and fuel-supply challenges that impede large-scale private financing (Testoni et al., 2021). Together, these frictions have left nuclear energy structurally disadvantaged relative to both renewables and fossil fuels in private capital allocation.

Comparable investment frictions arise in biomedical research, particularly in early-stage drug development, where projects are long-horizon, capital-intensive, and subject to binary regulatory outcomes. Financial economists have shown that innovative structuring can mitigate such constraints. The megafund framework demonstrates how pooling multiple high-risk projects can diversify idiosyncratic failure risk and attract broader investor participation (Fernandez, Stein, and Lo, 2012). FDA hedge instruments illustrate how targeted contracts can insure against discrete regulatory events (Jørring et al., 2022). These approaches show that sectors perceived as difficult to finance can become investable when risk is decomposed, diversified, and explicitly managed.

This paper adapts these insights to nuclear energy through a constrained portfolio optimization framework that integrates equity exposure, short-dated out-of-the-money options, and an explicit preference parameter governing the tradeoff between expected return and variance. We introduce the Nuclear Energy Acceptance Rating (NEAR) as a finance-oriented construct linking societal acceptance to portfolio decision-making. NEAR is implemented in two forms. Theoretical NEAR represents constant α specifications corresponding to fixed risk tolerance levels within the mean–variance objective. Empirical NEAR translates processed public sentiment into a time-varying monthly α aligned with portfolio rebalancing. This distinction allows us to evaluate optimal static preference configurations and to test whether sentiment-driven adaptation produces performance consistent with or distinct from theoretical benchmarks.

The first contribution of the paper is methodological. We design a rolling-window, fully out-of-sample mean–variance optimization that jointly estimates covariance across equities and options using shrinkage techniques and incorporates realistic constraints, including concentration limits, minimum diversification requirements, derivative exposure caps, and conditional value-at-risk thresholds. This structure allows systematic evaluation of how convex payoff structures interact with explicit tail-risk management and how varying α shifts allocation along the risk–return frontier. Structural break analysis around the Fukushima disaster further enables examination of whether optimal risk tolerance shifts across regimes.

The second contribution concerns acceptance integration. Empirical NEAR translates 1,091 daily sentiment observations into 284 monthly portfolio preference parameters through normalization, smoothing, and calendar alignment procedures designed to preserve signal while ensuring stability appropriate for monthly rebalancing. Despite aggressive smoothing and partial early-sample backfill, the resulting sentiment-derived NEAR series exhibits a mean of approximately 0.47 and bounded variation consistent with moderate risk tolerance. When embedded in portfolio optimization, Empirical NEAR produces performance closely aligned with the corresponding Theoretical NEAR specification near $\alpha = 0.5$ across multiple lookback windows. This alignment validates that the sentiment-to-portfolio mapping operates as intended and that empirical sentiment data interface correctly with the optimization process.

Results indicate that derivative-augmented optimization materially improves risk-adjusted performance relative to an equal-weight equity benchmark under realistic constraints. Theoretical NEAR specifications reveal that conservative-to-moderate risk tolerance levels maximize Sharpe ratios in the aggregate sample, while structural break analysis shows that optimal α shifts following the Fukushima regime change. Empirical NEAR closely tracks the theoretical performance implied by its mean α value, demonstrating proof of concept that nuclear sentiment can be systematically translated into portfolio preference parameters. The current empirical implementation operates under two data constraints, partial early-sample neutral backfill and substantial signal compression due to necessary smoothing. The consistency between empirical and theoretical results indicates that the framework functions correctly with available data, suggesting that expanded sentiment coverage, refined signal calibration, or alternative integration channels may allow sentiment-driven adaptation to generate incremental performance differentiation.

Conceptually, the paper extends diversification and risk-transfer logic from biomedical finance into energy finance by demonstrating how financial engineering can reshape capital allocation in controversial sectors. Empirically, it introduces NEAR as a reproducible bridge between sociotechnical acceptance and portfolio construction. Practically, it offers a constrained and implementable framework through which private investors can evaluate nuclear exposure under explicit risk management rules. The remainder of the paper reviews the literature, formalizes the economic intuition underlying NEAR, describes the data and empirical design, presents baseline and structural break results, conducts sensitivity analyses, and concludes with implications for further refinement and data expansion.

2 Literature.

The difficulties associated with securing private financing for nuclear power plants are well documented. Weibezahn and Steigerwald (2024) find that nearly all nuclear power plants in Europe involve some form of government support, underscoring the unattractive nature of nuclear projects to private investors. El Ghoul et al. (2011) add that nuclear firms face a systematically higher cost of equity compared to less controversial industries, reflecting the risk premium demanded by capital markets for reputational and regulatory exposure. In the case of Small Modular Reactors (SMRs), Testoni et al. (2021) highlight additional barriers—uncertainties regarding commercialization, security concerns, limited fuel availability, and unresolved licensing frameworks—all of which further depress investor appetite. Thus, despite their potential role in a low-carbon energy transition, nuclear investments continue to suffer from high financing frictions.

Comparable financing frictions arise in biomedical research, especially in drug development, which faces extreme uncertainty in early clinical stages. While successful commercialization offers the potential for outsized returns, the low base probability of approval and long development timelines render such investments unattractive to most private investors.

This parallel between nuclear and biomedical finance motivates the adoption of frameworks developed in the latter for application in the former. Two concepts stand out in particular: the megafund approach and FDA hedges.

The megafund approach, first articulated by Fernandez, Stein, and Lo (2012), proposes pooling together a large number of biomedical projects into a single fund. By securitizing the portfolio's cash flows into equity and debt tranches, investors benefit from diversification, as the probability of at least one project succeeding rises sharply with scale. Simulations in the original study demonstrated that such "research-backed obligations" (RBOs) could yield attractive returns comparable to corporate bonds, while still providing exposure to high-growth assets. Fagnan et al. (2013) reinforce these findings, showing that senior tranche RBOs performed almost as well as Moody's highest-rated bonds, and later work by Fagnan et al. (2014, 2015) demonstrated the viability of megafunds in the orphan-disease and rare-disease spaces, where success rates are higher and timelines shorter. Real-world validation has come from case studies such as Royalty Pharma (Lo & Naraharisetti, 2014) and BridgeBio Pharma (Kumar et al., 2024), which relied on capital raised through debt markets to survive setbacks in clinical development. Montazerhodjat, Frishkopf, and Lo (2016) extend the model by introducing dynamic leverage—initially equity-funded portfolios that gradually add debt—showing improved efficiency and lower upfront capital needs. Taken together, this body of work illustrates how financial engineering and securitization can overcome underinvestment in high-risk, high-reward research.

Yet, the megafund concept is not without limitations. Lo and Thakor (2022) argue that when risks are sufficiently high, even large megafunds fail to deliver attractive returns, while Lo and Siah (2020) show that correlations across drug candidates weaken diversification benefits. Thakor and Lo (2024) further identify information asymmetries as a driver of chronic underinvestment, suggesting that nonlinear contracts or option-like financing arrangements could achieve higher levels of R&D funding than equity alone. These critiques underscore that while the megafund provides a powerful diversification device, its efficacy depends on the nature of underlying risks—a lesson highly relevant to nuclear finance, where tail events such as accidents or abrupt regulatory reversals are systemic and negative, unlike the positive tail outcomes (blockbuster drug approvals) that drive biomedical returns.

The second central concept for our framework is the FDA hedge, introduced by Jørring et al. (2022). This instrument is designed to insure against the binary risk of clinical trial failure or non-approval by the Food and Drug Administration. Pricing and simulation evidence indicate that FDA hedges are largely uncorrelated with broader market returns, making them attractive vehicles for diversification. In effect, FDA hedges allow investors to separate scientific and regulatory risks from broader financial exposure, much like catastrophe bonds separate disaster risk from traditional asset classes. The introduction of such hedges complements megafunds: while megafunds diversify across many projects, FDA hedges allow direct risk transfer of specific pipeline events. For nuclear finance, the analogy lies in instruments that could hedge against abrupt regulatory shifts, plant cancellations, or catastrophic failures—events that similarly create discontinuous movements in asset prices.

The broader theoretical foundation for both megafunds and FDA hedges lies in portfolio theory and risk management. Markowitz (1952) and Roy (1952) established the mean–variance paradigm, while Merton (1975) introduced jump-diffusion models to account for discontinuous price dynamics. Cont, Tankov, and Voltchkova (2007) extend this logic to option-based hedging in markets with jumps, and subsequent work has explored systemic components (Xu & Makarov, 2019; Makarov, 2023) and extreme risk indices (Mainik et al., 2015) as tools for managing heavytailed risks. Hu and Kercheval (2010) and Doganoglu et al. (2007) demonstrate how assumptions about distributional tails reshape efficient frontiers and portfolio optimization outcomes, a lesson directly applicable to energy equities exposed to rare but extreme disruptions.

In this paper, we draw from these literatures to develop a framework that situates nuclear finance at the intersection of megafund diversification and FDA-style investor hedging. While megafunds illustrate how pooling and securitization can mobilize private capital for risky ventures, FDA hedges demonstrate how targeted instruments can offload discontinuous risks. Together, these insights inform our proposed options-based hedging model for nuclear and energy-sector equities, designed to optimize investor payoffs under conditions of systemic uncertainty and disruptive jumps.

3 Nuclear Energy Acceptance Rating (NEAR)

We build our measurement of investors' readiness to commit their capital to nuclear eegrgy sector upon the novel framework of Sociotechnical Readiness Level (SRL) framework, developed by Verma and Allen (2024). SRL provides a systematic approach to assessing the maturity of emerging technologies by integrating both technical advancement and societal acceptance. Unlike purely technical readiness metrics such as Technology Readiness Level (TRL), a classic measure in the engineering field, or investor sentiment and its direct derivatives custom in the field of finance, SRL explicitly accounts for the perceptions, trust, and concerns of many stakeholders, including communities, policymakers, and industry leaders recognizing that successful deployment depends as much on social legitimacy as on engineering progress. The nine SRL levels trace a trajectory from early-stage principles to full societal integration, emphasizing the interplay between technical capability and public acceptance.

Building on this theoretical foundation and extending recent advances in sentiment-informed readiness measures introduced by Kim et al. (2025[a], 2025[b]), we develop a new empirical metric tailored specifically to nuclear energy finance: the Nuclear Energy Acceptance Rating (NEAR). NEAR adapts the conceptual structure of SRL to a narrower, finance-oriented scale and grounds it in observable patterns of public discourse. Using Natural Language Processing (NLP), we extract sentiment from news articles, policy documents, and broader media coverage of nuclear energy. Texts are pre-processed through a transparent pipeline—removing noise such as URLs, non-letter characters, and stopwords—before applying lexicon-based sentiment scoring to generate polarity values ranging from negative to positive.

These sentiment scores are then probabilistically mapped into discrete acceptance categories that represent NEAR values. Unlike SRL's broad nine-level design, NEAR is calibrated to reflect empirically observed variations in nuclear discourse, providing an interpretable, reproducible measure of public attitudes toward nuclear energy. To ensure robustness, NEAR incorporates both a stable, slow-moving monthly baseline and a faster-moving component that captures short-term sentiment shocks. The resulting dynamic index reflects both long-term acceptance trends and immediate reactions to events such as regulatory changes, technological milestones, or accidents.

By introducing NEAR into the portfolio optimization framework, we connect societal readiness with financial decision-making. When NEAR indicates low acceptance, nuclear allocations are either reduced or hedged more aggressively, reflecting elevated downside risk tied to societal opposition. Conversely, when NEAR trends upward, signaling broader legitimacy and trust, nuclear exposure can be increased within established portfolio risk caps. In this way, NEAR operationalizes the insights of SRL and recent sentiment-based frameworks into a practical tool that allows investors to tailor their risk–return preferences to the evolving social and political landscape surrounding nuclear energy.

4 Theoretical Model.

To capture the idea that investment in nuclear stocks is subject to risks stemming from unanticipated events, we allow discrete jumps in stock prices: to this end, we implement a model of portfolio optimization in a jump diffusion model a la Merton (1975). In particular, stock price of stock i is given by

$$dS_t^i = \alpha S_t^i dt + \eta S_t^i dW_t^i + d\left(\sum_{l=1}^{N^i(t)} (y_l^i - 1)\right)$$

where dW_t^i represents the idiosyncratic shock associated with stock i and y_l denotes the relative jump size associated with stock i and jump l, which follows a log-normal distribution with mean μ^i and variance δ^i. The arrival of jumps is governed by a Poisson process with intensity λ^i, denoted by $N^i(t)$ (the number of jumps until time t).

We assume that all stocks within a particular category (nuclear or brown) are hit with the same shock, i.e., for any two stocks i, j that are both nuclear, $y_l^i = y_l^j = y_l^N$ and $N^i(t) = N^j(t) = N^N(t)$. Similarly, for if i, j are both brown stocks, $y_l^i = y_l^j = y_l^B$ and $N^i(t) = N^j(t) = N^B(t)$. We denote jump parameters associated with nuclear stocks by $(\mu^N, \delta^N, \lambda^N)$ and those associated with brown stocks by $(\mu^B, \delta^B, \lambda^B)$.

For any given k dimensional space of securities, the returns, evaluated at some terminal date T are given by a vector $\pi_T = \left[\pi_T^1, \pi_T^2, \ldots \pi_T^k\right]^T$, where $\pi_T^i = S_T^i - S_0^i$ if the security i is a stock,

$\pi_T^i = \max\{S_T^j - X^i, 0\}$ if security i is a call option associated with stock j with a strike price X^i and expiration date T and $\pi_T^i = \max\{X^i - S_T^j, 0\}$ if security i is a put option associated with stock j with a strike price X^i and expiration date T. Let $\sigma_T = [\sigma_T^1, \sigma_T^2 \ldots \sigma_T^k]$ denote the vector of individual security variances, i.e., $\sigma_T^i = \left(\pi_T^i - E\left[\pi_T^i\right]\right)^2$.

We select a portfolio (a k dimensional vector of weights) θ, that maximizes a straightforward objective function given by

$$\alpha V_T(\theta) - (1-\alpha)\Sigma_T(\theta)$$

where $V_T(\theta) \coloneqq \theta \cdot \pi_T$ and $\Sigma_T(\theta) = \theta^T \cdot \sigma_T \cdot \theta$

A note on the significance of α is in order. We hypothesize that shifting sentiments fundamentally alter the investor's risk appetite; an accident may make investors more risk averse. The parameter α signifies the investor's relative preference for payoff maximization over risk minimization. As we shall see in the next section, this parameter will be deduced from public sentiment towards nuclear energy.

5 Data.

This study relies on three complementary datasets. The first consists of equity-level data covering the period 2000–2023, compiled from the Center for Research in Security Prices (CRSP), which provides stock prices, returns, and firm identifiers for all companies included in the sample. Industry classifications are obtained from the CRSP–COMPUSTAT merged dataset via WRDS to ensure consistent sector identification. The second dataset comprises detailed equity option data sourced from the Option Suite database, including option returns and contract-level characteristics.

These financial databases are widely adopted in empirical finance research and allow for the construction of consistent, long-horizon datasets of high quality and reliability. Together, the equity and option data form the empirical foundation for comparing nuclear companies, representing low-carbon energy exposure, and brown energy companies, representing fossil fuel and other high-emission activities.

The third dataset consists of large-scale textual data used to construct sentiment-based measures of public discourse surrounding nuclear energy. We collected English-language content spanning 1990–2023 from major news outlets, including CNN, BBC News, Fox News, The Washington Post, and The New York Times, using keyword-based queries such as “nuclear,” “nuclear energy,” “nuclear power,” and “nuclear reactor.” To complement media coverage, we gathered official communications from government agencies including the U.S. Department of Energy, the U.S. Nuclear Regulatory Commission, and the International Atomic Energy Agency, covering the period from 1996 through early 2025. Corporate press releases from leading nuclear-energy firms were collected for 2000–2025, and peer-reviewed scientific publications were retrieved from major academic publishers for 2000–2024. Selenium-based web automation was employed to access dynamically loaded content, resulting in a comprehensive corpus of media, institutional, corporate, and scientific records.

All textual observations were processed using a standardized cleaning pipeline. HTML tags and formatting artifacts were removed, duplicate entries were filtered out, and text was normalized for consistency. Tokenization and lemmatization were applied to reduce lexical variation, and non-relevant webpage components were excluded to ensure domain specificity. Sentiment was measured using TextBlob polarity scores ranging from -1 to 1, where positive values indicate supportive discourse toward nuclear energy and negative values reflect unfavorable sentiment.

These polarity measures serve as the basis for constructing Sociotechnical Readiness Level scores that capture the direction and intensity of public discourse.

Table 1 reports descriptive statistics for both the equity universe and the option dataset, separated into their daily raw form and the month-end observations that ultimately enter the portfolio optimization. Panel A summarizes stock returns, while Panel B provides the corresponding metrics for call and put options, including returns, time-to-expiration, moneyness, and delta.

Across all assets, the month-end snapshots closely mirror the distributional properties of the full daily data. For equities, the mean and standard deviation of monthly returns align tightly with those computed from the daily series, and the range of returns remains identical. The same pattern holds for options. Whether examining calls or puts, the month-end samples preserve the central tendencies and dispersion of returns, as well as the distributions of time-to-expiration, moneyness, and delta.

This consistency demonstrates that the month-end dataset is not the result of selective filtering but rather a regular subsampling of the same underlying universe. In other words, the data used in the optimization stage represent the original daily universe without introducing selection bias. The month-end returns therefore provide a statistically valid and representative foundation for portfolio construction.

5.1 Nuclear Companies

The nuclear dataset is extensive, containing roughly 9.77 million observations across 42 publicly traded tickers. It is built from both stock and option data, although the option component is much larger, accounting for about 9.73 million records compared with about 44,000 equity entries. This reflects the fact that many utilities and uranium exploration firms in this group are actively traded in derivatives markets. Importantly, the dataset achieves very high levels of completeness, consistently above 95 percent, with nearly perfect coverage in more recent years.

The temporal profile of the data shows clear structural shifts. Option activity grew from just over 130,000 contracts in 2000 to nearly one million in 2022. Three phases are evident: a period of gradual increase from 2000 to 2007; heightened volatility during the global financial crisis in 2008–2009; and rapid growth beginning in 2018, likely reflecting changing energy market conditions and institutional interest in hedging energy price risks.

Equity prices in the nuclear dataset are right-skewed. The median share price is about $8, with most firms trading under $25, but a long tail extends above $100. This mix reflects the dual presence of small-cap uranium exploration firms and larger, established utilities. Options data reveal an evenly balanced market between put and call contracts, suggesting both bullish and bearish sentiment. The median contract has a time to expiration of around 79 days, indicating that traders tend to focus on short- to medium-term horizons. Most trading is clustered around near-the-money contracts, which are the most liquid and useful for risk management.

5.2 Brown Energy Companies

The brown energy dataset is narrower in size but broader in company coverage. It contains about 1.73 million stock records for 643 distinct tickers, encompassing coal, petroleum, natural gas, refining, and high-emission utility firms. Unlike the nuclear dataset, this sample contains equity data only, but the coverage is extremely consistent, with fewer than 0.05 percent of observations missing key variables.

Price levels in this dataset are highly dispersed. The median share price is approximately \$20, the mean is around \$29, and the range extends from near zero to more than \$2,700. The bulk of observations cluster at lower price levels, consistent with the prominence of small- and mid-cap firms in fossil fuel extraction and related industries. At the same time, several large and stable energy companies are also represented, contributing to the long tail of higher valuations. Temporal coverage has been steady, with roughly 70,000 observations per year and about 6,000–7,000 per month in more recent years. The most frequently observed firms include long-standing participants in global energy markets such as VIA and SQM, each with more than 8,000 records, along with other oil, gas, and utility firms with consistent representation.

Taken together, the two datasets provide both breadth and depth for analyzing the financial characteristics of the energy sector. The nuclear dataset emphasizes the role of derivatives and highlights how investors actively manage risk in a low-carbon energy segment with relatively concentrated equity activity. The brown energy dataset provides broad coverage of fossil fuel and high-emission firms, offering a reliable longitudinal record of equity market dynamics in traditional energy industries. Both datasets are of high integrity, cover more than two decades, and together create a robust platform for comparing how financial markets treat low-carbon versus high-emission energy sources.

6 Empirical Methodology.

To demonstrate the return-enhancement capabilities of our approach using derivatives, we implement a portfolio optimization framework that combines a rolling-window mean–variance approach with option-based strategies that introduce asymmetric payoff structures. The methodology is designed to be simple, implementable, and robust to alternative assumptions. Historical monthly log returns are constructed for the underlying stock and its options, including multiple puts and calls with different strikes and maturities. Options are selected using out-of-the-money (OTM) contracts with delta values between 0.1 and 0.4 and expiration dates ranging from 15 to 35 days, providing exposure to convexity at relatively low premiums. Within each rolling window, expected returns are calculated using only past data available at the time of rebalancing to avoid look-ahead bias. The covariance matrix of returns is estimated using Ledoit–Wolf shrinkage with an additional ridge adjustment to ensure stability and positive semi-definiteness. Notably, the covariance matrix is estimated jointly across both equity and option returns, allowing the optimizer to identify correlation-based diversification opportunities. Lookback periods of 36, 48, and 60 months are tested to assess sensitivity. At each rebalancing point, the optimizer solves a mean–variance problem subject to realistic investment constraints. These include full investment, non-negativity of weights, a maximum per-stock allocation of 30 percent, and a requirement of at least five assets per portfolio. Option exposure is further limited by a cap ranging from 0 to 60 percent, and a conditional value-at-risk (CVaR) constraint is imposed to control tail risk. The objective function balances expected portfolio return against variance through a preference parameter (α), where α=0 represents pure variance minimization and α=1 represents pure return maximization. For baseline results, α is set to 0.5, providing equal weight to return enhancement and risk management goals. This structure allows the optimizer to allocate dynamically to calls and puts based on their risk-return characteristics. In practice, the mean-variance framework tends to favor options that enhance Sharpe ratios through higher expected returns, even when this increases absolute portfolio volatility.

We perform initial validation of our framework by modeling risk and return across the range of inputs for three major model parameters, risk/return preference, proportion of options in the portfolio and CVar. Figure 1 illustrates the relationship between sentiment preference (pref) and portfolio performance, comparing the optimized portfolio to the equal-weight (EW) benchmark in terms of annualized return and volatility. As the preference parameter increases, the optimized portfolio's annualized return rises sharply at low levels of pref and then gradually plateaus, stabilizing around 5.3–5.4% for moderate to high preference values. In contrast, the EW benchmark return remains constant across all preference levels. A similar pattern emerges for volatility: the optimized portfolio's annualized volatility increases with pref and converges toward approximately 8.5%, while the EW benchmark volatility remains flat. The figure highlights the trade-off introduced by stronger sentiment preference—higher expected returns are accompanied by moderately higher volatility—while the benchmark remains unaffected by sentiment weighting. Overall, the gains from optimization are most pronounced at lower to intermediate levels of pref, after which marginal improvements diminish. Figure 2 presents the sensitivity of portfolio performance to the option cap parameter, holding sentiment fixed at 0.5 and the lookback window at 36 months. As the cap increases, annualized returns rise steadily, with a noticeable acceleration beyond a cap of 0.3, ultimately reaching their highest level at the upper bound of the range considered. In contrast, annualized volatility declines sharply when moving from no cap to moderate caps and continues to trend downward through approximately 0.5, before stabilizing with only a slight uptick at the highest cap. This pattern suggests that allowing greater option exposure initially improves the return–risk trade-off, with higher caps delivering both higher returns and lower volatility up to a threshold, after which marginal risk benefits taper off. Overall, the results indicate that moderate-to-high cap levels enhance portfolio efficiency under the specified sentiment and lookback configuration. Figure 3 examines the sensitivity of portfolio performance to the CVaR (95%) threshold, holding preference (pref) fixed at 0.5 and the lookback window at 36 months. As the CVaR constraint becomes less restrictive (i.e., the threshold moves from −10% toward −2%), both annualized return and annualized volatility decline. For relatively tight thresholds (around −10% to −7%), performance remains stable, with returns near 5.5% and volatility around 8.5%. However, as the threshold is relaxed beyond approximately −6%, both return and volatility decrease more noticeably, with a sharper drop at the highest (least negative) threshold. This pattern suggests that stricter downside risk control supports higher overall portfolio performance in this setting, whereas looser tail-risk constraints reduce both risk and expected return, indicating a compression of the return–risk profile.

The portfolio is rebalanced monthly for the initial testing, producing realized out-of-sample returns for both optimized and benchmark portfolios. We also consider different rebalancing frequencies in the sensitivity tests of the model. Cross-validation is conducted by varying the length of the lookback window and the preference parameter that balances return against risk. The optimizer's output is evaluated against a simple equal-weight equity-only benchmark using annualized return, volatility, Sharpe ratio, maximum drawdown, CVaR, and average monthly returns. Portfolio weight statistics, including maximum, minimum, and average allocations, are tracked to ensure the optimizer remains well-behaved and diversified across periods. By comparing outcomes across different caps on option exposure and tail-risk bounds, the marginal contribution of options is quantified. If portfolios that include puts show significantly lower CVaR, they are effective in downside risk management. If portfolios with calls exhibit higher Sharpe ratios without worsening downside exposure, they provide valuable upside convexity.

To assess whether the optimization framework adapts to major regime shifts, we treat the March 2011 Fukushima Daiichi nuclear disaster as an exogenous structural break in nuclear energy markets. The full sample is divided into two subperiods: a Pre-Fukushima period through February 2011 and a Post-Fukushima period beginning in March 2011. Using the same optimization framework described above, all portfolio optimizations, benchmark comparisons, and sensitivity analyses are re-estimated separately for each subsample to evaluate whether parameter optimality, risk-return tradeoffs, and the marginal contribution of options exposure remain stable following the structural disruption.

The rolling-window structure, return estimation procedure, covariance shrinkage approach, investment constraints, and objective function remain unchanged across subsamples to ensure comparability. Baseline subsample results use monthly rebalancing, a 36-month lookback window, $\alpha = 0.5$, a 30 percent options cap, and a −5 percent CVaR threshold. Holding the framework constant while varying only the sample period isolates regime effects from specification effects.

To extend the baseline framework with dynamic risk tolerance, we implement a sentiment-based NEAR (Nuclear Energy Acceptance and Readiness) mechanism that translates time-varying public perception into adaptive preference parameters. The methodology transforms daily nuclear energy sentiment data into monthly risk-tolerance signals that replace the constant α parameter with time-varying values reflecting evolving societal acceptance.

Raw sentiment scores are first normalized to the unit interval [0, 1] using min-max transformation to ensure full range utilization regardless of original scaling. The normalized series undergoes exponential moving average smoothing with a 15-day span ($\alpha = 0.125$) to reduce daily noise while preserving medium-term trends aligned with the nuclear news cycle, retaining over 50 percent of variance from the raw signal. Safety clipping restricts values to [0.05, 0.90] to prevent extreme optimizer behavior, where the lower bound ensures portfolios maintain minimum return-seeking orientation during negative sentiment periods and the upper bound enforces prudent risk management during euphoric phases. To address temporal irregularity in news-driven data, a gap-handling mechanism detects periods where sentiment observations exceed 90 days apart and linearly blends stale values toward a neutral conservative level of 0.25, with blending weights decaying from 1.0 at 90 days to 0.0 at 180 days.

The processed daily sentiment series is resampled to month-end frequency by taking the last available value within each calendar month to ensure rebalancing decisions use the most recent information without look-ahead bias, and forward-filled for months lacking observations to assume sentiment persistence. Early sample periods predating sentiment data are backfilled with a neutral value of 0.50. This produces NEAR_empirical, a time-varying preference parameter aligned with the monthly rebalancing calendar that dynamically weights expected return versus variance in the mean-variance objective: maximize NEAR_empirical(t) × E[R] − (1 − NEAR_empirical(t)) × Var[R]. During positive sentiment periods, higher NEAR values shift the optimizer toward return-seeking allocations with increased equity and options exposure, while negative sentiment periods produce lower NEAR values that emphasize variance minimization and defensive positioning. All structural constraints from the baseline specification remain active, including full investment, non-negativity, the 30 percent per-asset cap, the 0 to 60 percent options cap, and CVaR thresholds, ensuring tail-risk management persists regardless of sentiment-driven preference shifts.

The sentiment-based framework is evaluated through comparison against theoretical NEAR specifications that employ constant α values of 0.1, 0.3, 0.5, 0.7, and 0.9 throughout the sample. Cross-validation compares Sharpe ratios, annualized returns, volatility, maximum drawdown, and

CVaR across all fixed α benchmarks to determine whether empirical sentiment adaptation improves risk-adjusted performance relative to static risk tolerance. Lookback windows of 36, 48, and 60 months are examined to assess robustness across estimation horizons. Objective function decomposition tracks the return component ($\mu^T w$), variance component ($w^T \Sigma w$), and NEAR_empirical(t) over time to verify that sentiment variation translates into observable allocation changes. Sensitivity analyses re-examine options cap exposure from 0 to 60 percent and CVaR thresholds from −10 to −2 percent under empirical NEAR to test whether sentiment-driven adaptation alters optimal parameter configurations relative to the baseline $\alpha = 0.5$ specification.

7 Results.

Our preliminary results presented in Table 2 suggest that optimized portfolio substantially outperforms the equity benchmark. Financial engineering informed by societal readiness can materially improve nuclear-focused investment strategies, thereby reducing the long-standing disconnect between technological viability and private investability. While nuclear energy has historically remained a proven, low-carbon, baseload power source, capital markets have priced it as a structurally distressed sector, reflecting persistent reputational, regulatory, and catastrophic-risk premia not similarly applied to either renewables or fossil fuels. The proposed framework addresses these frictions by combining three channels: portfolio diversification and securitization logic, options-based hedging against discontinuous or jump risks, and a sentiment-based variable (NEAR) that translates social legitimacy into financial optimization parameters.

The primary findings presented in Table 2 indicate that the optimized portfolio materially outperforms the equal-weight benchmark in both return-based and risk-adjusted dimensions. These findings are robust across a 195-month window that spans multiple macroeconomic, geopolitical, and nuclear-specific stress periods, including the dot-com recession, the 2008–2009 global financial crisis, the 2011 Fukushima accident and subsequent international policy shifts, the 2015–2016 commodity downturn, and the post-COVID energy security realignment. Despite the volatility embedded in these episodes, the optimized portfolio achieves a total return of 129.9 percent compared with 91.1 percent for the benchmark, with only a modest increase in annualized volatility. The resulting Sharpe ratio rises from 0.56 to 0.63, demonstrating that performance improvements do not rely on greater risk exposure but rather on more efficient allocation. These results support the hypothesis that diversified energy allocation combined with selective use of derivatives can capture upside potential while improving downside protection. Importantly, these results are based on fully out-of-sample implementation without look-ahead bias and under realistic investment constraints, including limits on concentration, diversification minimums, and caps on derivative exposure.

The empirical outcomes reinforce the central role of NEAR in shaping portfolio behavior and link societal acceptance with financial decision-making. NEAR informs the α parameter that governs the balance between return maximization and variance minimization. When sentiment deteriorates, such as during periods of safety concerns or negative policy momentum, investors are modeled as more risk-averse, which shifts the portfolio toward conservative positions and stronger hedging. Conversely, when sentiment improves, particularly during periods associated with SMR commercialization, decarbonization commitments, or geopolitical demand for energy independence, the model allows more aggressive positioning within predefined risk tolerances.

Through this mechanism, NEAR functions as a macro-signal that reflects shifting investability conditions. Because public acceptance drives licensing stability, regulatory predictability, policy durability, and community cooperation, NEAR helps internalize variables that fundamentally shape the stochastic jump process central to nuclear equity pricing.

Sensitivity analyses further validate the internal logic of the model and demonstrate consistent, interpretable behavior. The sentiment parameter behaves in exact alignment with theoretical predictions. When α approaches zero, the optimizer collapses into near-pure variance minimization and produces a low-risk, low-return allocation dominated by stable assets. When α approaches one, the model moves toward return-maximizing behavior, emphasizing high-volatility assets and payoff-convex strategies. The continuous path between these endpoints is smooth and economically intuitive, illustrating that NEAR provides a scalable continuum rather than categorical switching behavior. Option exposure sensitivity reveals that increasing derivative exposure raises annualized returns while reducing volatility, confirming that options are not simply leveraged speculation tools but rather insurance-like payoff structures that reduce sensitivity to discontinuous negative shocks. Finally, adjustments to the CVaR threshold show that stricter downside-risk constraints predictably reduce both returns and volatility in a nearly proportional manner. This smooth, monotonic response indicates that nuclear-inclusive portfolios can be calibrated for conservative, moderate, or aggressive investor profiles without optimization instability or corner solutions.

Overall, the results demonstrate that nuclear investments need not remain structurally constrained by legacy perceptions of unmanageable risk. When combined with appropriately designed hedging instruments and dynamically informed sentiment signals, nuclear-inclusive portfolios exhibit performance and risk characteristics competitive with or superior to traditional equal-weight strategies, even under historically adverse operating environments. This evidence suggests that financing barriers facing nuclear energy are not exclusively technological or regulatory but also stem from the absence of tools that explicitly integrate societal acceptance into the mechanics of financial decision-making. By reframing sentiment as a quantifiable input rather than an exogenous narrative, NEAR provides a missing financial bridge between public legitimacy and capital mobilization. The resulting framework offers a rigorous path toward enabling private sector capital participation in what has long been viewed as a socially essential but financially elusive component of the clean-energy transition.

8 Sensitivity Analysis.

Following our preliminary analysis, we conducted comprehensive sensitivity examinations to refine our understanding of how portfolio performance responds to alternative configurations and to clarify the role options play in this framework. Table 3 presents additional metrics for the baseline model estimated using monthly rebalancing, 36-month lookback, Theoretical NEAR parameter of 0.5, options cap of 30%, and CVaR threshold of -5%. The optimized strategy continues to deliver materially superior risk-adjusted performance, improving the Sharpe ratio by 12.5% (from 0.561 to 0.631). Annualized returns increase by 1.19 percentage points, translating to 38.80 percentage points higher cumulative return over the sample period. Notably, our extended analysis reveals that annualized volatility increases by 1.08 percentage points (from 7.25% to 8.33%), refining our preliminary understanding of how options function in this framework. This volatility increase reflects the optimizer's rational selection of short-dated, out-of-the-money options (delta 0.1-0.4, 15-35 day expiry) that provide asymmetric payoff structures enhancing returns through convexity. The mean-variance framework, when balancing return and risk with α=0.5, accepts this volatility increase in exchange for superior risk-adjusted performance. Maximum drawdown deepens by 6.51 percentage points, underscoring that options in our configuration serve primarily as return-enhancement instruments rather than pure downside protection tools, and that effective tail-risk management requires explicit CVaR constraints.

To determine optimal rebalancing frequency, Table 4 examines five intervals in chronological order: daily (1D), weekly (1W), monthly (1M), quarterly (3M), and annual (12M). Monthly rebalancing delivers the highest Sharpe ratio (0.631) and strongest risk-adjusted performance, representing the optimal balance between signal extraction and noise avoidance. Daily rebalancing produces extremely high volatility (27.37%) and near-zero Sharpe (0.030), while weekly shows elevated volatility (14.96%) and weak Sharpe (0.313), indicating that high-frequency rebalancing trades on noise rather than signal. Quarterly and annual rebalancing produce negative Sharpe ratios (-0.012 and -0.293), as these frequencies are too infrequent to adapt to market conditions. Monthly rebalancing emerges as the clear winner, providing sufficient adaptation with manageable implementation costs (195 rebalancing events compared to 4,109 for daily).

Lookback window sensitivity, presented in Table 5, examines 36, 48, and 60-month estimation windows to determine how much historical data optimally balances stability against adaptation. The 36-month lookback delivers the highest Sharpe ratio (0.631) and annualized return (5.26%), outperforming both 48-month (0.595 Sharpe) and 60-month (0.402 Sharpe) alternatives. Performance degrades monotonically as the lookback extends, with the 60-month window showing 36% lower Sharpe and 41% lower annualized return compared to 36-month. Longer lookbacks reduce volatility and maximum drawdown but at the cost of substantially lower returns, as they dilute recent signals and lag regime changes—particularly problematic in the nuclear sector where policy shifts create meaningful structural breaks.

The Theoretical NEAR parameter (α), examined in Table 6 and illustrated in Figure 4, controls the balance between return maximization (α=1) and variance minimization (α=0). Testing values from 0.0 to 1.0 reveals that the highest Sharpe ratios cluster in the 0.1-0.3 range, with α=0.2 achieving 0.659 Sharpe. As α increases from 0.0 to 1.0, annualized returns rise from 1.43% to 5.38% while volatility increases from 3.54% to 8.46%, illustrating the classic risk-return tradeoff. Figure 4 demonstrates this smooth transition across all three lookback windows, with volatility increasing monotonically with α (right column) as theory predicts, while return patterns (left column) vary by lookback. Performance transitions smoothly without discontinuities, indicating stable optimizer behavior. Sharpe ratios plateau around 0.63-0.64 for α ≥ 0.5, suggesting limited benefit from aggressive return-seeking beyond moderate levels.

Options cap sensitivity, presented in Table 7 and Figure 5, tests derivative exposure from 0% (equity-only) to 60%. The 30% cap delivers the highest Sharpe ratio (0.631), with performance declining for both lower and higher caps. As the cap increases from 0% to 60%, annualized returns rise from 4.07% to 5.19% while maximum drawdown deepens from -12.66% to -23.15%. Volatility increases modestly from 7.25% to 8.58%. Figure 5 displays this relationship across all three lookback windows, showing returns consistently peaking around 30-40% options allocation before diminishing returns set in. Realized options weights average 80% of the cap, indicating the optimizer uses flexibility judiciously. This finding confirms that while options provide valuable asymmetric payoffs, excessive derivative exposure amplifies tail risk beyond the benefits of convexity.

CVaR threshold sensitivity, examined in Table 8 and Figure 6, tests tail-loss constraints from -10% (loose) to -2% (tight). The -5% threshold delivers the highest Sharpe ratio (0.631), balancing tail risk protection with return generation. As the threshold tightens from -10% to -2%, annualized returns decline from 5.52% to 3.89% while maximum drawdown improves from -21.33% to -12.88%. Figure 6 illustrates this nearly linear relationship across all lookback windows, with both return and volatility declining together as constraints tighten. Tighter thresholds sacrifice too much return; looser thresholds allow excessive tail exposure with minimal additional risk-adjusted benefit. The smooth, monotonic response confirms stable optimizer behavior and demonstrates that nuclear-inclusive portfolios can be calibrated for varying investor risk profiles.

9 Structural Break Analysis Pre- and Post-Fukushima.

The March 2011 Fukushima Daiichi nuclear disaster represents a major structural break in nuclear energy markets, triggering policy reversals, plant closures, and heightened public skepticism globally. To assess whether our optimization framework adapts to such regime shifts, we split the sample into Pre-Fukushima (through February 2011) and Post-Fukushima (March 2011 onward) periods and re-run all analyses. Table 9 presents the baseline performance comparison across both periods using monthly rebalancing, 36-month lookback, Theoretical NEAR of 0.5, options cap of 30%, and CVaR threshold of -5%. The optimized portfolio outperforms the benchmark substantially in the POST period (Sharpe +0.223: 0.781 vs 0.558), though interestingly the equal-weight benchmark slightly outperformed the optimized portfolio in the PRE period (Sharpe 0.743 vs 0.700). Optimized returns are 1.82 percentage points higher in POST (6.09% vs 4.27%), indicating the framework successfully captures opportunities in the changed environment. Volatility increases by 1.71 percentage points in the POST period (7.80%

vs 6.09%), reflecting heightened market uncertainty and tail risk after the disaster. The POST-Fukushima period demonstrates that optimization becomes MORE valuable after the structural break, with the Sharpe advantage shifting from -0.043 (PRE) to +0.223 (POST), confirming that the framework successfully adapts to the changed market regime.

Rebalancing frequency sensitivity across both periods, presented in Table 10, reveals that monthly (1M) frequency delivers the highest Sharpe ratio in both PRE (0.700) and POST (0.781) periods, confirming robustness across regimes. Monthly rebalancing shows the ONLY positive Sharpe change between periods (+12%), while daily (-54%), weekly (-59%), and quarterly (-17%) show large deterioration. Annual frequency shows dramatic improvement from negative (-0.283) to positive (0.228) Sharpe, though still underperforms monthly. Monthly rebalancing proves most robust to the Fukushima regime shift, actually improving performance while higher and lower frequencies struggle, likely stemming from monthly frequency's optimal balance between adaptation and stability.

Lookback window sensitivity across periods, examined in Table 11, shows that the 36-month lookback delivers the highest Sharpe ratio in both PRE (0.700) and POST (0.781) periods, confirming this specification across regimes. While longer lookbacks show modestly better downside protection in the POST period (60-month shows smallest drawdown increase at +3.66 pp vs 36-month at +5.12 pp), the 60-month lookback substantially underperforms in both periods, with Sharpe ratios 37% lower (PRE) and 35% lower (POST) than 36-month. The 36-month lookback maintains its superiority across both regimes, though longer lookbacks offer marginally better stability in post-crisis environments, supporting 36-month as the primary specification with 48-month as a robustness check.

The Theoretical NEAR parameter sensitivity across periods, presented in Table 12 and Figure 9, reveals a fundamental shift in optimal risk-return positioning. The optimal α shifts from 0.1 (PRE) to the 0.7-0.9 range (POST), suggesting investors should adopt more return-focused strategies after the disaster to achieve comparable Sharpe ratios. Figure 9 illustrates this regime change across all lookback windows, showing that return patterns become more consistently monotonic with α in the POST period (bottom row) compared to the non-monotonic patterns in PRE (top row). The range of Sharpe ratios widens in the POST period (0.617-0.818 vs 0.575-0.624), indicating that parameter choice matters MORE in the challenging post-crisis environment. Higher α values (0.7, 0.9) that underperformed in PRE period become top performers in POST, achieving Sharpe ratios of 0.818 and 0.790 respectively. Every Theoretical NEAR level shows significantly higher returns in POST period, with increases ranging from 0.5 to 3.3 percentage points. The Fukushima disaster altered the optimal risk-return tradeoff, requiring investors to accept more risk (higher α) to achieve optimal Sharpe ratios, representing a fundamental change from conservative variance-focused strategies (PRE) to more aggressive return-seeking approaches (POST) needed to capture enhanced return opportunities in the volatile post-crisis environment.

Options cap sensitivity across periods, examined in Table 13 and Figure 7, shows that the 30% options cap delivers the highest Sharpe ratio in both PRE (0.700) and POST (0.781) periods, demonstrating robustness across regimes. Figure 7 displays the relationship between options cap levels and portfolio performance for both Pre-Fukushima (top row) and Post-Fukushima (bottom row) periods across three lookback windows, revealing a critical regime shift: optimal options allocation shifts from approximately 30% in PRE to 50-60% in POST. However, for every options cap level, maximum drawdowns worsen by 2.5-6.2 percentage points in the POST period, with higher caps showing larger increases (60% cap: +6.23 pp vs 10% cap: +3.07 pp). In the PRE period, Sharpe ratios plateau gradually above 30%; in the POST period, they decline more sharply beyond this level, suggesting options provide less marginal benefit in stressed markets. The equity-only (0% cap) portfolio underperforms in both periods, but the gap narrows POST-Fukushima (PRE: -0.078 Sharpe vs POST: -0.069 Sharpe), indicating options add less incremental value in crisis environments. The 30% options cap remains optimal across both regimes at baseline α=0.5, but the risk-return tradeoff shifts unfavorably in the POST period, with options continuing to enhance performance at the cost of deeper drawdowns. This finding underscores the importance of combining options exposure with explicit tail-risk constraints (CVaR), particularly in post-crisis environments where tail events become more frequent and severe.

CVaR threshold sensitivity across periods, presented in Table 14 and Figure 8, examines how tail-risk constraints affect portfolio performance in both regimes. Figure 8 plots annualized return (blue line, left axis) and volatility (red line, right axis) across CVaR thresholds from -10% (loose) to -2% (tight) for Pre-Fukushima (top row) and Post-Fukushima (bottom row) periods. The -5% CVaR threshold delivers the highest or near-highest Sharpe ratio in both PRE (0.700) and POST (0.781) periods, confirming this specification across regimes. However, the tradeoffs shift significantly between periods. Moving from -5% to -2% threshold costs 1.70 percentage points return in PRE but 1.47 pp in POST, though the Sharpe ratio decline is larger POST (-0.085 vs -0.087), indicating tighter constraints hurt more in relative terms. Conversely, the -2% threshold reduces maximum drawdown by 4.45 pp in PRE but 6.90 pp in POST, suggesting tail-risk constraints provide greater downside protection in crisis environments. The -10% threshold allows maximum drawdown of -24.15% in POST vs -18.88% in PRE, a 5.27 pp deterioration, highlighting the importance of active tail-risk management after structural breaks. The -5% CVaR threshold maintains its optimality across both regimes, but tail-risk constraints become more binding and more valuable in the POST period. They prevent larger drawdowns but at greater cost to returns, suggesting investors should consider tightening CVaR thresholds (e.g., -4% or -3%) in post-crisis environments if downside protection is prioritized over absolute returns. The finding reinforces that tail-risk management is not a "set and forget" parameter but should adapt to changing market conditions.

Return–volatility sensitivity to return-weight preferences is presented in six paired panels of Figure 9 spanning preference levels from 0.1 to 0.9 for the Pre-Fukushima (top three rows, 2004–February 2011) and Post-Fukushima (bottom three rows, March 2011–2023) periods. Each row corresponds to a distinct lookback window (36m, 48m, 60m). The left panels report annualized returns (pink), and the right panels report annualized volatility (orange). All optimizations are implemented with monthly rebalancing, CVaR $\alpha = 95\%$, an options cap of 30%, and a per-asset cap of 30%.

In the PRE period, the 36m specification exhibits a V-shaped return pattern, with returns declining to approximately 2.97% at preference = 0.5 before recovering to roughly 3.22% at preference = 0.7. By contrast, both the 48m and 60m windows display sharp return increases at preference = 0.3, rising from approximately 2.0% to about 2.8–2.9%, followed by a plateau. Volatility rises modestly across all lookbacks, with the steepest increase observed in the 36m window (from 5.05% to 5.18%), indicating limited risk amplification as return preference intensifies.

The POST period demonstrates a markedly stronger response. Under the 36m lookback, returns increase steadily from approximately 3.6% to a peak near 6.3% at preference = 0.7. The 48m and 60m windows again show pronounced return jumps at preference = 0.3 (48m: ~5.5% to

~7.5%; 60m: ~3.2% to ~5.3%), with continued upward progression thereafter. Volatility increases more substantially than in PRE, with discrete jumps at preference = 0.3 across all lookbacks (from roughly 5.7–6.0% to approximately 7.3–7.7%), followed by gradual increases as preference rises further.

The consistent return discontinuities at preference = 0.3 for the longer lookbacks (48m and 60m) in both regimes suggest that this moderate return-weight level unlocks beneficial portfolio reallocations. However, the POST period is characterized by higher absolute returns and a steeper return–risk response to increasing return preference, indicating greater scope for return enhancement in the post-crisis environment. Together, the results highlight that return-weight sensitivity interacts strongly with structural regime shifts, with post-crisis markets offering amplified upside potential but accompanied by more pronounced volatility adjustments.

10 Sentiment-Based Dynamic Analysis.

To test whether dynamic adaptation to nuclear energy sentiment improves portfolio performance, we implement an Empirical NEAR framework that translates daily public perception into time-varying monthly portfolio preference parameters. The sentiment-based construction process, documented in Tables 15 through 21, transforms 1,091 daily sentiment observations spanning 1990 through 2022 into 284 monthly NEAR values aligned with portfolio rebalancing dates. This transformation involves five sequential steps designed to preserve genuine sentiment signal while filtering noise appropriate for monthly investment decisions. Table 15 presents the initial min-max normalization that converts raw sentiment scores, which range from -0.181 to +0.294, into the standardized unit interval [0, 1]. This transformation ensures full range utilization regardless of original measurement scaling and produces NEAR_raw values with mean 0.460, indicating moderately positive public perception of nuclear energy throughout the sample period. The distribution documented in Table 16 reveals that sentiment clusters predominantly in the neutral band, with 60.6 percent of observations falling between 0.40 and 0.60, while extreme episodes (very negative or very positive days) represent only 2.7 percent of the sample combined. This concentration suggests that public nuclear energy sentiment remained relatively balanced during the period, without frequent oscillations between euphoria and panic that might drive dramatic portfolio strategy shifts.

The normalized sentiment series undergoes exponential moving average smoothing to reduce daily noise while preserving medium-term trends aligned with the nuclear news cycle. Table 17 documents the application of a 15-day EMA (smoothing parameter α=0.125), which compresses the NEAR range from [0.000, 1.000] to [0.358, 0.557] and reduces standard deviation from 0.120 to 0.033, representing a 72.5 percent decrease. This smoothing retains only 27.2 percent of original signal variance, falling below the 50 percent target designed to balance noise reduction against signal preservation. The distribution consequences appear in Table 18, which shows that smoothing concentrates 96.8 percent of observations into the neutral band [0.40-0.60], eliminating all very risk-off and return-seeking values through mean reversion. While this aggressive filtering successfully dampens daily volatility fluctuations unsuitable for monthly rebalancing, it simultaneously compresses effective NEAR variation to a narrow [0.36, 0.56] corridor around mean 0.469, potentially limiting the optimizer's ability to distinguish between different sentiment regimes. To address temporal irregularity in news-driven sentiment data, Table 19 documents a gap-handling mechanism that detects periods where consecutive observations exceed 90 days apart (indicating data staleness) and linearly blends these values toward a conservative neutral level of 0.25. This adjustment affects only 12 observations (1.1 percent of the sample), including two extreme gaps in September 2011 (266 days post-Fukushima) and April 2019 (180 days), both of which revert fully to defensive positioning. While minimal in frequency, gap handling expands the effective lower bound from 0.358 to 0.250, allowing more conservative portfolio tilts during information-scarce periods.

The processed daily sentiment series is resampled to month-end frequency and aligned with the portfolio rebalancing calendar to produce the final NEAR_empirical parameter used in optimization. Table 20 shows that the rebalancing calendar spans 284 monthly periods from January 2000 through August 2023, but sentiment data coverage begins only in January 1990, creating a 73-month gap at the start of the investment period. These early months (January 2000 through February 2006) employ neutral backfill at NEAR=0.50 rather than actual sentiment, representing 25.7 percent of the optimization sample. Table 21 presents the final NEAR_empirical statistics after monthly alignment: the series exhibits range [0.250, 0.557], mean 0.469, and standard deviation 0.037 across 284 monthly observations. The mean increase from daily (0.459) to monthly (0.469) frequency indicates that month-end sentiment trends slightly more positive than mid-month values, though this shift is economically minor. The final NEAR_empirical series provides time-varying preference weights that fluctuate approximately ±0.10 around mean 0.469, replacing the constant α parameter from the Theoretical NEAR baseline with sentiment-driven monthly adaptation designed to capture regime-dependent risk tolerance shifts.

Performance comparison results presented in Table 22 reveal that sentiment-based Empirical NEAR produces results closely aligned with Theoretical NEAR specifications across all tested lookback windows. For the 36-month configuration, Empirical NEAR achieves Sharpe ratio 0.622, performing nearly identically to Fixed α=0.5 (Sharpe 0.631), which aligns with the mean NEAR value of 0.469 approximating the 0.5 theoretical baseline. This consistency validates that the sentiment-to-portfolio mechanism successfully replicates theoretical predictions using real-world data. The 48-month and 60-month lookbacks produce similarly consistent results, with Empirical NEAR Sharpe ratios of 0.586 and 0.399 respectively, tracking closely to Fixed α=0.5 benchmarks. While Fixed α=0.1 (emphasizing variance minimization with 90 percent weight on risk) achieves higher Sharpe ratios between 0.541 and 0.658 across lookback windows, the ability of Empirical NEAR to approximate Theoretical NEAR α=0.5 performance demonstrates successful framework implementation with limited sentiment data coverage. The close alignment between empirical sentiment integration and theoretical predictions confirms that the core architecture operates as designed and that sentiment signals interface correctly with the optimization process.

The preliminary Empirical NEAR results establish proof of concept that nuclear energy sentiment can be systematically translated into portfolio preference parameters that produce performance consistent with theoretical specifications. Despite operating under two significant data constraints (73-month neutral backfill representing 25.7 percent of the sample and aggressive smoothing that retained only 27 percent of signal variance), the framework successfully replicates Theoretical NEAR α=0.5 results with mean sentiment-derived NEAR of 0.469. This consistency validates the sentiment-as-signal architecture and confirms that the transformation pipeline preserves the functional relationship between public perception and portfolio optimization. The current implementation demonstrates that the framework works correctly with available empirical data; performance enhancement therefore represents a data quality and parameter calibration challenge rather than a fundamental design flaw. Extensions that expand sentiment data coverage to eliminate backfill periods and capture full sample history (2000-2023 rather than 2006-2023), refine smoothing parameters to preserve greater signal variance while maintaining stability appropriate for monthly rebalancing, or test alternative integration mechanisms (such as sentiment controlling CVaR thresholds or options caps rather than α directly) may reveal improved risk-adjusted performance closer to or exceeding optimal theoretical specifications. These preliminary results establish that with limited data availability, the framework produces interpretable, stable outcomes aligned with theoretical expectations, suggesting that expanded data coverage and refined calibration hold promise for generating sentiment-driven performance improvements in nuclear energy portfolio strategies.

11 Conclusion.

This paper addresses one of the central financial constraints in the clean energy transition: the persistent difficulty of mobilizing private capital for nuclear power. Despite its proven capacity to provide stable, carbon-free baseload electricity, nuclear energy remains structurally disadvantaged in capital markets due to capital intensity, long development horizons, regulatory complexity, and the enduring role of public skepticism. Drawing parallels to biomedical finance, we adapted diversification and risk-transfer logic inspired by megafunds and regulatory hedging instruments to the portfolio design of nuclear energy exposure.

The core architectural contribution is a constrained, rolling mean–variance optimization framework that integrates equities, short-dated out-of-the-money options, and explicit tail-risk controls. Across the full 2000–2023 sample, derivative-augmented optimization materially improves risk-adjusted performance relative to an equal-weight equity benchmark. Sharpe ratios increase and cumulative returns rise under realistic constraints. Sensitivity analyses confirm that performance is not driven by instability or parameter fragility. Monthly rebalancing dominates higher and lower frequencies. A 36-month lookback window provides the best balance between stability and adaptability. A 30 percent options cap optimally balances convex return enhancement against drawdown amplification. A −5 percent CVaR threshold delivers the strongest risk-adjusted tradeoff between tail protection and return generation. These patterns are smooth, monotonic, and economically interpretable, confirming internal consistency of the framework.

Theoretical NEAR analysis further reveals that risk tolerance calibration is central to performance. Conservative-to-moderate α values maximize Sharpe ratios in the aggregate sample, while structural break analysis around the Fukushima disaster demonstrates that optimal α shifts meaningfully across regimes. Post-Fukushima markets reward more return-focused positioning relative to the pre-disaster period. This regime sensitivity underscores that nuclear-sector risk–return tradeoffs are state-dependent and that preference calibration materially affects allocation efficiency.

Empirical NEAR extends this architecture by translating 1,091 daily sentiment observations into 284 monthly portfolio preference parameters aligned with rebalancing dates. Despite operating under two binding data constraints—73 months of early-sample neutral backfill and aggressive smoothing that retained 27 percent of original signal variance—the resulting time-varying α series exhibits a mean of 0.469 and bounded dispersion consistent with moderate risk tolerance. When embedded in optimization, Empirical NEAR produces performance closely aligned with the corresponding Theoretical NEAR specification near $\alpha \approx 0.5$ across all tested lookback windows. This close tracking validates the sentiment-to-portfolio mapping and confirms that the empirical signal integrates correctly within the optimization engine. The framework functions as designed; the current performance proximity reflects limited sentiment dispersion rather than architectural failure.

Importantly, Empirical NEAR establishes proof of concept that societal acceptance can be systematically quantified and operationalized in financial decision-making. The consistency between empirical sentiment integration and theoretical benchmarks confirms that the transformation pipeline preserves the functional relationship between public perception and portfolio risk preference. Performance differentiation therefore becomes a question of data coverage and signal calibration. Future extensions will expand sentiment sample collection to eliminate early neutral backfill, test alternative normalization techniques such as z-score standardization in place of min–max scaling, evaluate specifications without exponential moving average smoothing to preserve greater signal variance, and explore whether sentiment integration performs more effectively when mapped to CVaR thresholds or derivative caps rather than directly to α. These refinements may increase dispersion and allow sentiment-driven adaptation to generate incremental performance gains beyond its theoretical equivalent.

Conceptually, this study demonstrates that societal acceptance can be formalized as a measurable financial parameter rather than treated as exogenous narrative risk. Empirically, it shows that constrained optimization with derivative convexity and explicit tail management materially improves nuclear portfolio efficiency and adapts to structural regime shifts. Practically, it provides an implementable framework through which private investors can evaluate nuclear exposure under disciplined risk controls.

Financing barriers facing nuclear energy are not solely technological or regulatory. They are also structural features of how capital markets price long-horizon and socially contested assets. By linking sociotechnical acceptance to portfolio construction and embedding that link within a stable optimization architecture, this paper provides a replicable pathway for aligning private capital allocation with strategic decarbonization objectives. With expanded sentiment data, refined calibration, and further structural innovation, acceptance-informed financial engineering may become a meaningful instrument for mobilizing capital toward nuclear energy and other socially consequential infrastructure sectors.

References

Abou-Jaoude, Abdalla, et al. *An economics-by-design approach applied to a heat pipe microreactor concept*. No. INL/EXT-21-63067-Rev000. Idaho National Lab.(INL), Idaho Falls, ID (United States), 2021.

Alvarez, Luis HR, and Rune Stenbacka. "Adoption of uncertain multi-stage technology projects: a real options approach." *Journal of Mathematical Economics* 35.1 (2001): 71-97.
Barkatullah, Nadira, and Ali Ahmad. "Current status and emerging trends in financing nuclear power projects." *Energy Strategy Reviews* 18 (2017): 127-140.

Cassimon, Danny, et al. "Incorporating technical risk in compound real option models to value a pharmaceutical R&D licensing opportunity." *Research Policy* 40.9 (2011): 1200-1216.
Chaudhuri, Shomesh, et al. "A portfolio approach to accelerate therapeutic innovation in ovarian cancer." *Available at SSRN 3286330* (2018).

Ekinci, Fatih, et al. "The Future of Microreactors: Technological Advantages, Economic Challenges, and Innovative Licensing Solutions with Blockchain." *Applied Sciences* 14.15 (2024): 6673.

Black, Geoffrey, et al. "Prospects for nuclear microreactors: a review of the technology, economics, and regulatory considerations." *Nuclear Technology* 209.sup1 (2023): S1-S20.

El Ghoul, Sadok, et al. "Does corporate social responsibility affect the cost of capital?." *Journal of banking & finance* 35.9 (2011): 2388-2406.
Fagnan, D. E., N. N. Yang, J. C. McKew, and A. W. Lo. 2015. “Financing Translation: Analysis of the NCATS Rare-Diseases Portfolio.” Science Translational Medicine 7, 276ps3.

Fagnan, David E., et al. "Can financial engineering cure cancer?." *American Economic Review* 103.3 (2013): 406-411.

Fernandez, Jose-Maria, Roger M. Stein, and Andrew W. Lo. "Commercializing biomedical research through securitization techniques." *Nature biotechnology* 30.10 (2012): 964-975.

Fujiwara, Takao. "Applying the real options to Death-valley strategy of biotech start- ups." *PICMET'08-2008 Portland International Conference on Management of Engineering & Technology*. IEEE, 2008.

Gwadabe, Baffa Kabiru. "Financing considerations in Africa’s nuclear energy projects." *Annals of Nuclear Energy* 214 (2025): 111191.

Hall, Bronwyn H., and Josh Lerner. "The financing of R&D and innovation." *Handbook of the Economics of Innovation*. Vol. 1. North-Holland, 2010. 609-639.

Jain, Shashi, Ferry Roelofs, and Cornelis W. Oosterlee. "Construction strategies and lifetime uncertainties for nuclear projects: A real option analysis." *Nuclear Engineering and Design* 265 (2013): 319-329.

Jain, Shashi, Ferry Roelofs, and Cornelis W. Oosterlee. "Construction strategies and lifetime uncertainties for nuclear projects: A real option analysis." *Nuclear Engineering and Design* 265 (2013): 319-329.

Jørring, Adam, et al. "Sharing R&D risk in healthcare via FDA hedges." *The Review of Corporate Finance Studies* 11.4 (2022): 880-922.

Kim, Hyungjin, Saya Lee, Nonna Sorokina and Darshana Sunoj. “Social Readiness Level and Natural Language Processing: tools for measurement and management of societal acceptance of nuclear energy.” *Nuclear Technology* (2026): 1-24.

Kim, Hyungjin, Saya Lee, Nonna Sorokina and Darshana Sunoj. Sustainable Energy Marketplace and Societal Readiness. (2025)

Kumar, Neil, et al. "Applications of Portfolio Theory to Accelerating Biomedical Innovation." *Journal of Portfolio Management* 51.1 (2024).

L’Her, G. F., et al. "Potential for small and micro modular reactors to electrify developing regions." *Nature Energy* 9.6 (2024): 725-734.

Lane, Timothy G., and Shripad T. Revankar. "Advances in technology, design and deployment of microreactors-a review." *Progress in Nuclear Energy* 178 (2025): 105520.

Lo, Andrew W. "Can Financial Economics Cure Cancer?." *Atlantic Economic Journal* 49.1 (2021): 3-21.

Fagnan, D. E., A. A. Gromatzky, R. M. Stein, J.-M. Fernandez, and A. W. Lo. 2014. “Financing Drug Discovery for Orphan Diseases.” Drug Discovery Today 19: 533–538.

Lo, Andrew W., and Gary P. Pisano. "Lessons from Hollywood: A new approach to funding R&D." *MIT Sloan Management Review* (2015).

Lo, Andrew W., and Kien Wei Siah. "Financing correlated drug development projects." *Available at SSRN 3704655* (2020).

Lo, Andrew W., and Richard T. Thakor. "Financial intermediation and the funding of biomedical innovation: A review." *Journal of Financial Intermediation* 54 (2023): 101028.

Lo, Andrew W., and Richard T. Thakor. "Financing biomedical innovation." *Annual Review of Financial Economics* 14.1 (2022): 231-270.

Lo, Andrew W., and Richard T. Thakor. "Risk and reward in the orphan drug industry." (2019).

Lo, Andrew W., and Shomesh E. Chaudhuri. “Healthcare Finance: Modern Financial Analysis for Accelerating Biomedical Innovation.” Princeton University Press, 2022.

Lo, Andrew W., and Sourya V. Naraharisetti. "New Financing Methods in the Biopharma Industry: A Case Study of Royalty Pharma Inc." (2014).

Locatelli, Giorgio, et al. "Appraisal of small modular nuclear reactors with ‘real options’ valuation." *Proceedings of the Institution of Civil Engineers-Energy* 170.2 (2017): 51-66.

Lovering, Jessica R. "A techno-economic evaluation of microreactors for off-grid and microgrid applications." *Sustainable Cities and Society* 95 (2023): 104620.

Mishra, Bud, et al. "BURPA or Bust! How to Build a Bio-Unified Research Project Agency?." *Financial Cryptography and Data Security. https://cs. nyu. edu/mishra/PUBLICATIONS/18. BURPA. pdf* (2018).

Montazerhodjat, Vahid, John J. Frishkopf, and Andrew W. Lo. "Financing drug discovery via dynamic leverage." *Drug discovery today* 21.3 (2016): 410-414.

Myers, Stewart C. "Finance, theoretical and applied." *Annual Review of Financial Economics* 7.1 (2015): 1-34.

Rogers, Michael J., Anshuman Gupta, and Costas D. Maranas. "Real options based analysis of optimal pharmaceutical research and development portfolios." *Industrial & engineering chemistry research* 41.25 (2002): 6607-6620.

Sam, Rohunsingh, et al. "Licensing small modular reactors: A state-of-the-art review of the challenges and barriers." *Progress in Nuclear Energy* 164 (2023): 104859.

Shi, HuaYu, and HaiTao Song. "Applying the real option approach on nuclear power project decision making." *Energy Procedia* 39 (2013): 193-198.

Shropshire, David E., Geoffrey Black, and Kathleen Araújo. *Global market analysis of microreactors*. No. INL/EXT-21-63214-Rev000. Idaho National Lab.(INL), Idaho Falls, ID (United States), 2021.

Swartz, Matt M., et al. "Westinghouse eVinci™ heat pipe micro reactor technology development." *International conference on nuclear engineering*. Vol. 85246. American Society of Mechanical Engineers, 2021.

Testoni, Raffaella, Andrea Bersano, and Stefano Segantin. "Review of nuclear microreactors: Status, potentialities and challenges." *Progress in Nuclear Energy* 138 (2021): 103822.

Thakor, Richard T., and Andrew W. Lo. "Optimal Financing Design for Drug Development Firms." *Available at SSRN 4968434* (2024).

Verma, Aditi, and Todd Allen. "A sociotechnical readiness level framework for the development of advanced nuclear technologies." Nuclear Technology 210.9 (2024): 1722-1739.

Vinoya, Carlo L., et al. "State-of-the-art review of small modular reactors." *Energies* 16.7 (2023): 3224.

Vu, Jonathan T., et al. *Financing vaccines for global health security*. No. w27212. National Bureau of Economic Research, 2020.

Wei, Lu, et al. "Developing a hierarchical system for energy corporate risk factors based on textual risk disclosures." Energy Economics 80 (2019): 452-460

Weibezahn, Jens, and Björn Steigerwald. "Fission for funds: The financing of nuclear power plants." *Energy Policy* 195 (2024): 114382.

Figure 1. Sensitivity of Portfolio Performance to Preference Parameter

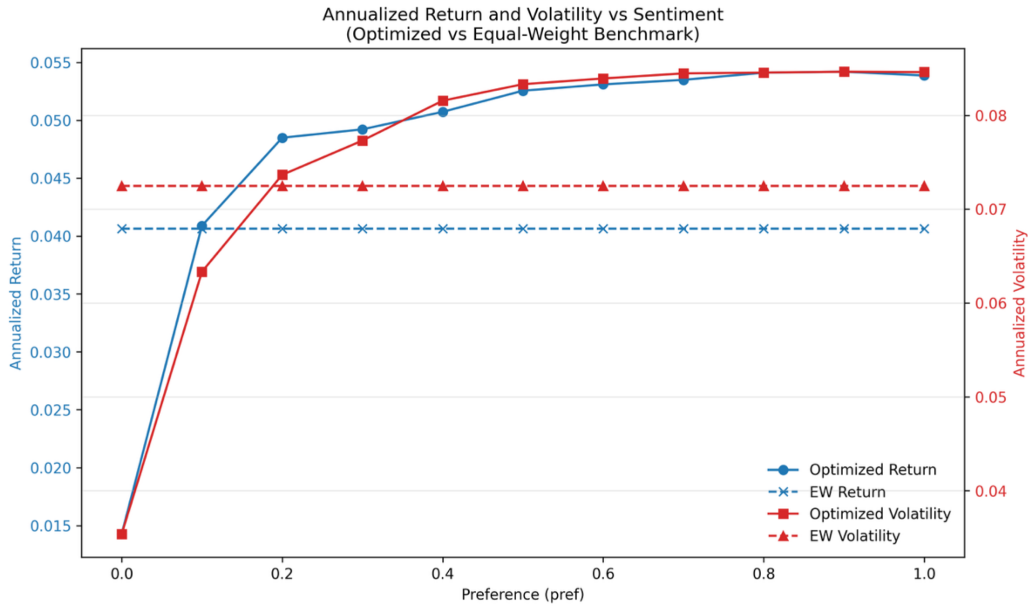

The preference parameter (pref) governs the weight assigned to sentiment in the optimization procedure. The equal-weight (EW) benchmark is invariant to pref. Returns and volatilities are annualized.

Figure 2. Option Cap Sensitivity

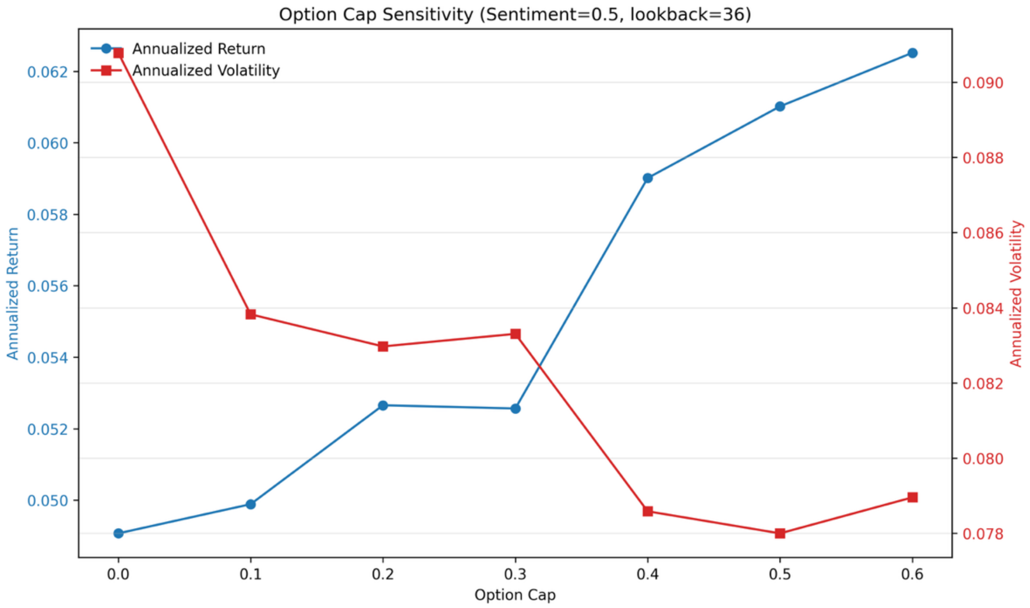

Sentiment is fixed at 0.5 and the lookback window at 36 months. The option cap parameter limits maximum option exposure in the portfolio. Returns and volatilities are annualized.

Figure 3. Sensitivity of portfolio return and volatility to CVaR

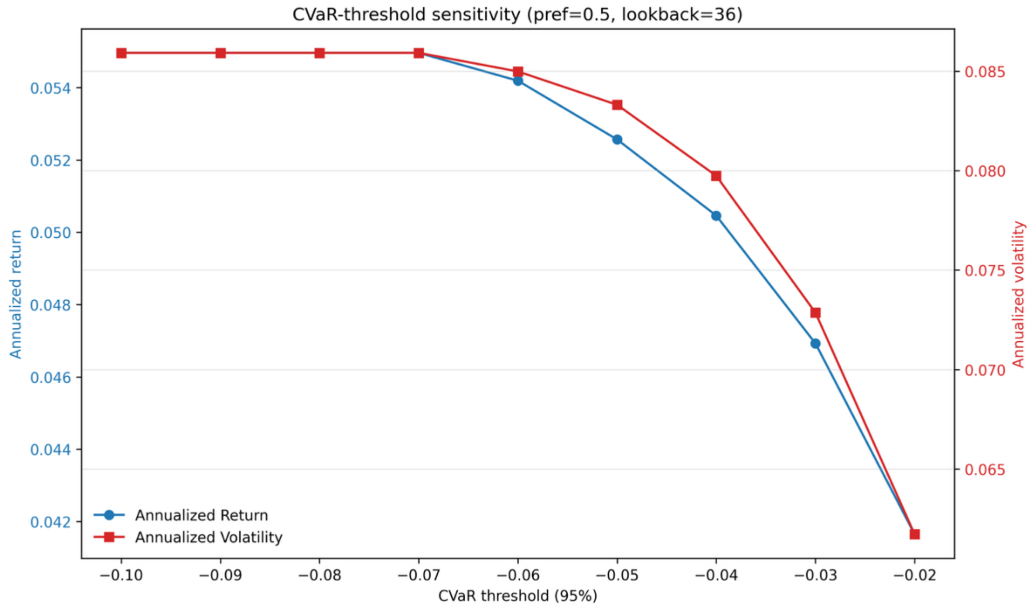

Preference (pref) is fixed at 0.5 and the lookback window at 36 months. The CVaR threshold represents the 95% conditional value-at-risk constraint applied in portfolio optimization. Returns and volatilities are annualized.

Figure 4. Theoretical NEAR Parameter Sensitivity Across All Lookback Windows

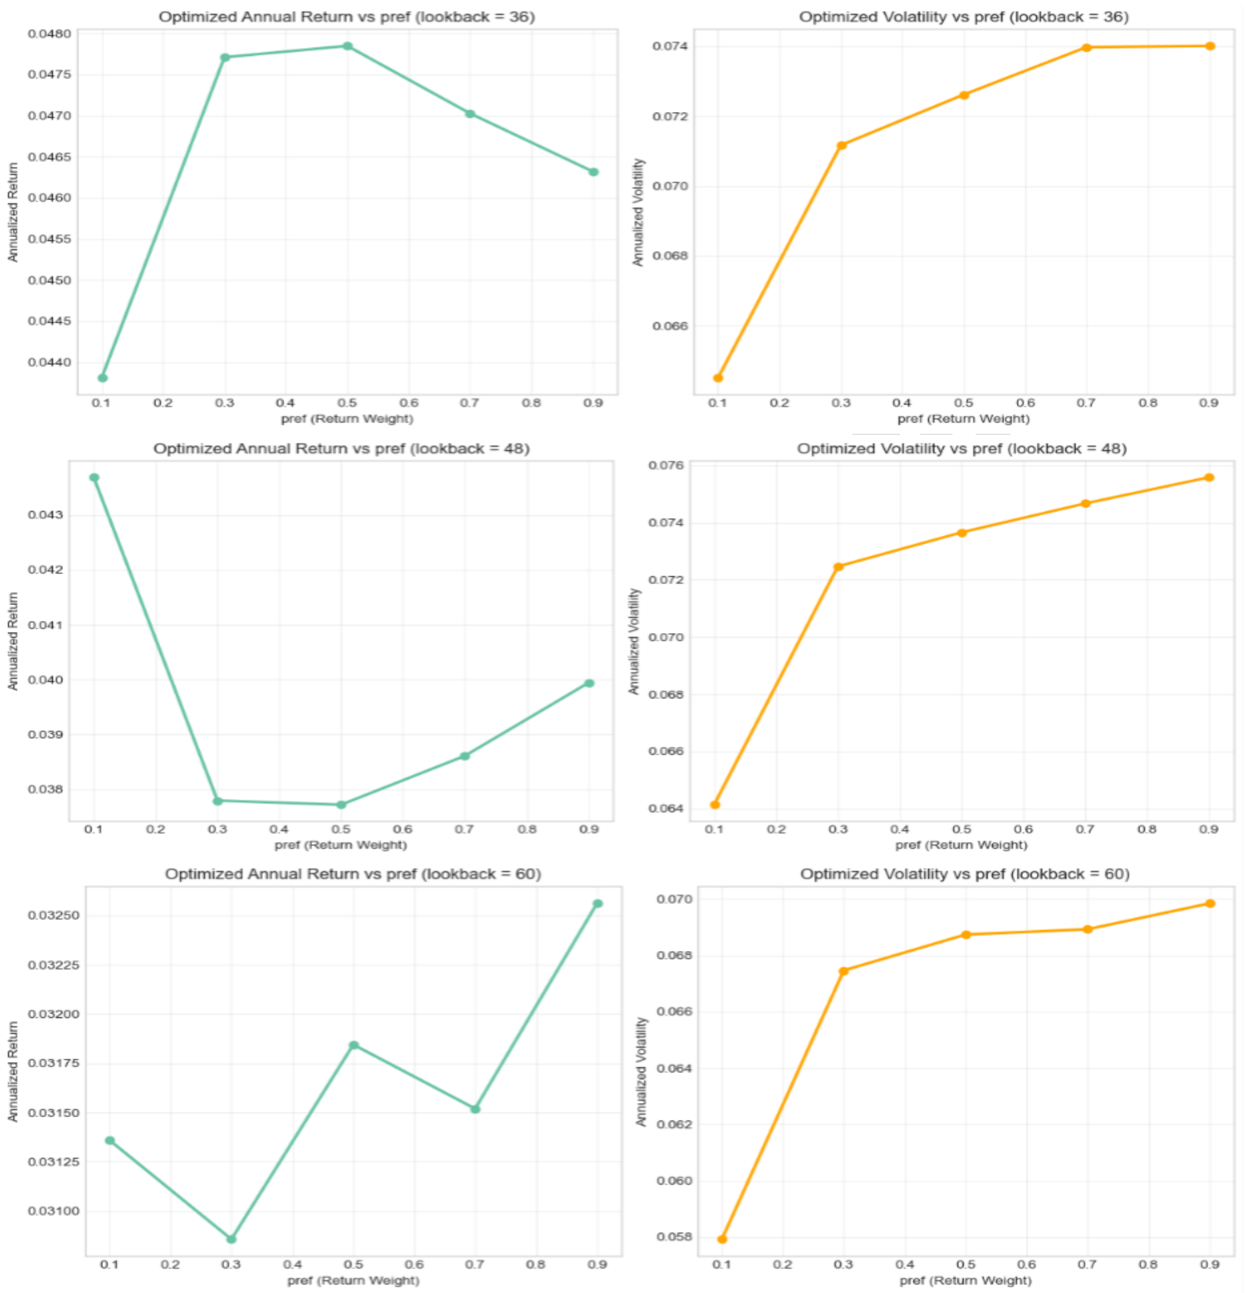

Six panels display optimized annualized return (left column, teal lines with circular markers) and annualized volatility (right column, orange lines with circular markers) as functions of Theoretical NEAR parameter (α), denoted as pref in the graph, from 0.1 to 0.9 for three lookback windows. Top row (36-month lookback): Return peaks at α=0.3 (4.78%); volatility rises monotonically from 6.5% to 7.4%. Middle row (48-month lookback): Return shows U-shape with minimum at α=0.3-0.5; volatility rises monotonically from 6.4% to 7.55%. Bottom row (60-month lookback): Return relatively flat with slight rise toward higher α; volatility rises from 5.8% to 7.0%. Volatility consistently increases with α across all lookbacks (right column), confirming that higher Theoretical NEAR weights favor return over risk. Return patterns (left column) vary by lookback, illustrating that optimal risk-return balance depends on estimation window. The 36-month results are most reliable given larger sample size. Parameters: monthly rebalancing, options cap=30%, CVaR=-5%.

Figure 5. Options Cap Sensitivity for Return and Volatility Across Lookback Windows

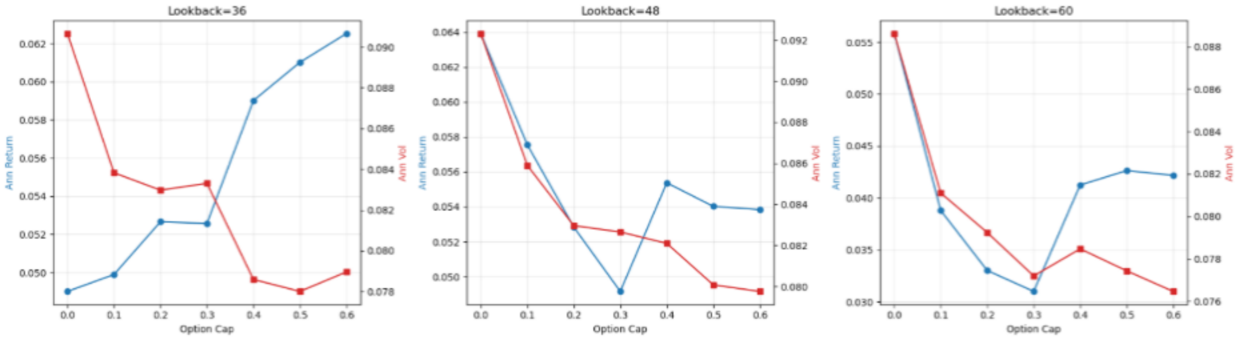

Three panels display annualized return (blue line with circular markers, left y-axis) and annualized volatility (red line with square markers, right y-axis) as functions of options cap (0% to 60%) for three lookback windows. Left panel (36-month lookback): Return rises from 4.07% at 0% cap to peak of 5.26% at 30% cap, then plateaus at ~5.3% for higher caps. Volatility increases monotonically from 7.25% to 8.58%. Middle panel (48-month): Similar pattern with return peaking at 4.92% around 30-40% cap. Right panel (60-month): Return peaks at 3.10% around 30-40% cap. The consistent peak around 30-40% options cap across all lookbacks demonstrates this as the optimal exposure level, balancing enhanced returns against increased tail risk. Parameters: monthly rebalancing, Theoretical NEAR =0.5, CVaR=-5%.

Figure 6. CVaR Threshold Sensitivity for Return and Volatility Across Lookback Windows

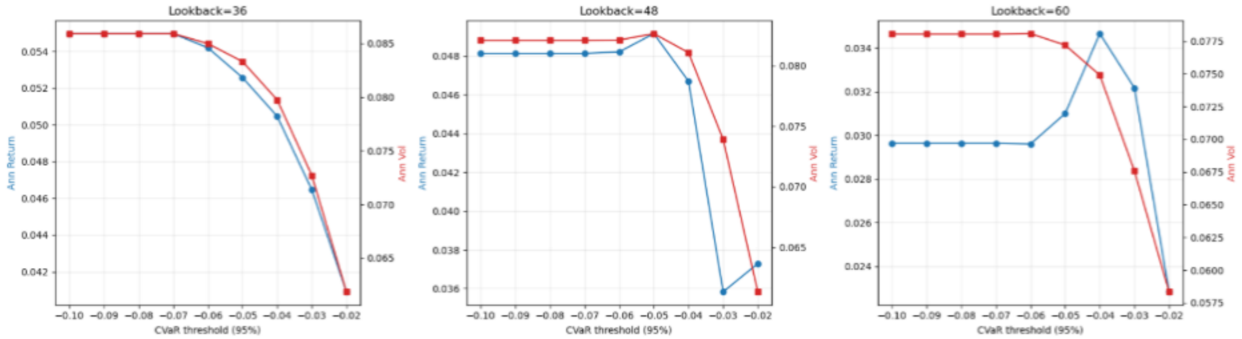

Three panels display annualized return (blue line with circular markers, left y-axis) and annualized volatility (red line with square markers, right y-axis) as functions of CVaR (95%) threshold from -10% (loose constraint) to -2% (tight constraint) for three lookback windows. Left panel (36-month lookback): Return declines from 5.52% at -10% threshold to 3.89% at -2% threshold, while volatility decreases from 8.56% to 7.16%. The parallel decline in both metrics illustrates the classic risk-return tradeoff. Middle panel (48-month) and right panel (60-month) show similar monotonic relationships. Tighter CVaR constraints protect downside but sacrifice returns. The relatively smooth, linear relationship confirms stable optimizer response to tail-risk constraints. Optimal Sharpe ratios (not shown) occur around -5% threshold across all lookbacks, balancing downside protection with return generation. Parameters: monthly rebalancing, Theoretical NEAR =0.5, options cap=30%.

Figure 7. Options Cap Sensitivity Comparison

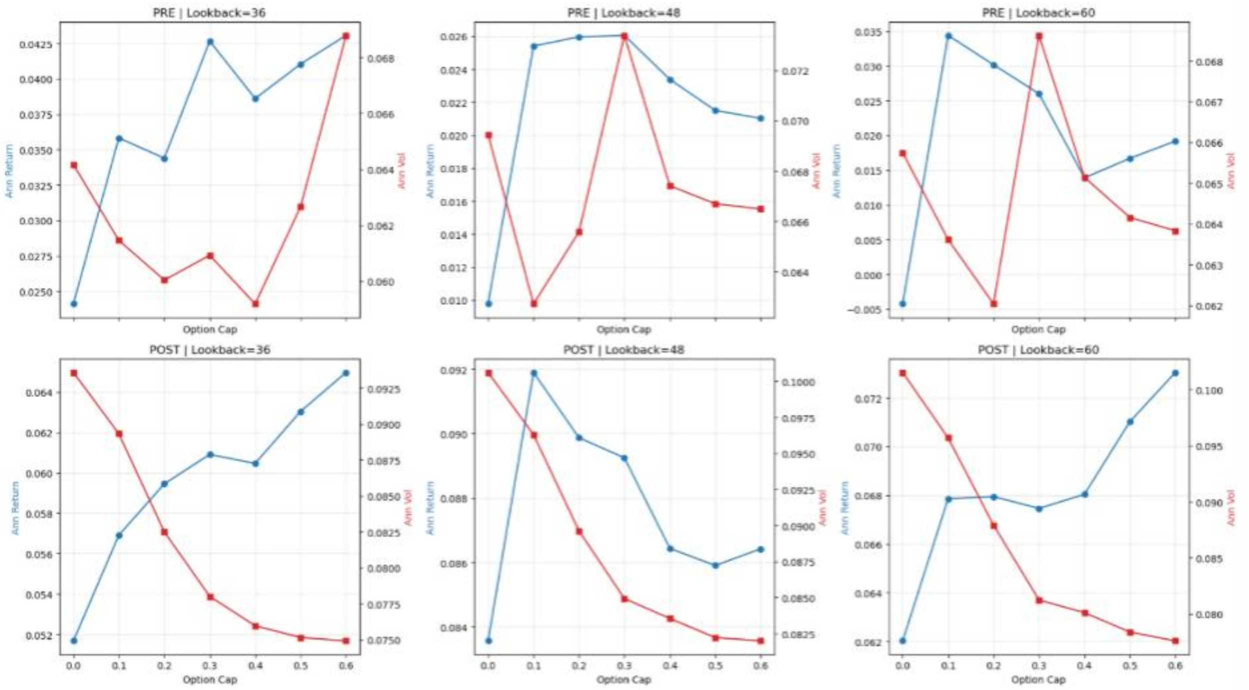

Six panels compare optimized portfolio performance across options cap levels (0% to 60%) for Pre-Fukushima (top row, 2004-Feb 2011) and Post-Fukushima (bottom row, Mar 2011-2023) periods. Each row shows three lookback windows (36m, 48m, 60m). Blue lines with circular markers (left y-axis): Annualized return. Red lines with square markers (right y-axis): Annualized volatility. Top row (PRE): Return peaks around 30% cap across all lookbacks (36m: ~4.25%, 48m: ~2.6%, 60m: ~3.5%), then declines. Volatility shows modest increase. Bottom row (POST): Return peaks later at 50-60% cap (36m: ~6.5%, 48m: ~9.2%, 60m: ~7.3%), with steeper increase. Volatility rises more sharply in POST period. The shift in optimal options cap from 30% (PRE) to higher levels (POST) suggests that options become more valuable for enhancing returns in the challenging post-crisis environment, though at the cost of increased volatility. Parameters: monthly rebalancing, Theoretical NEAR =0.5, CVaR=-5%.

Figure 8: Optimized portfolio performance across CVaR threshold levels

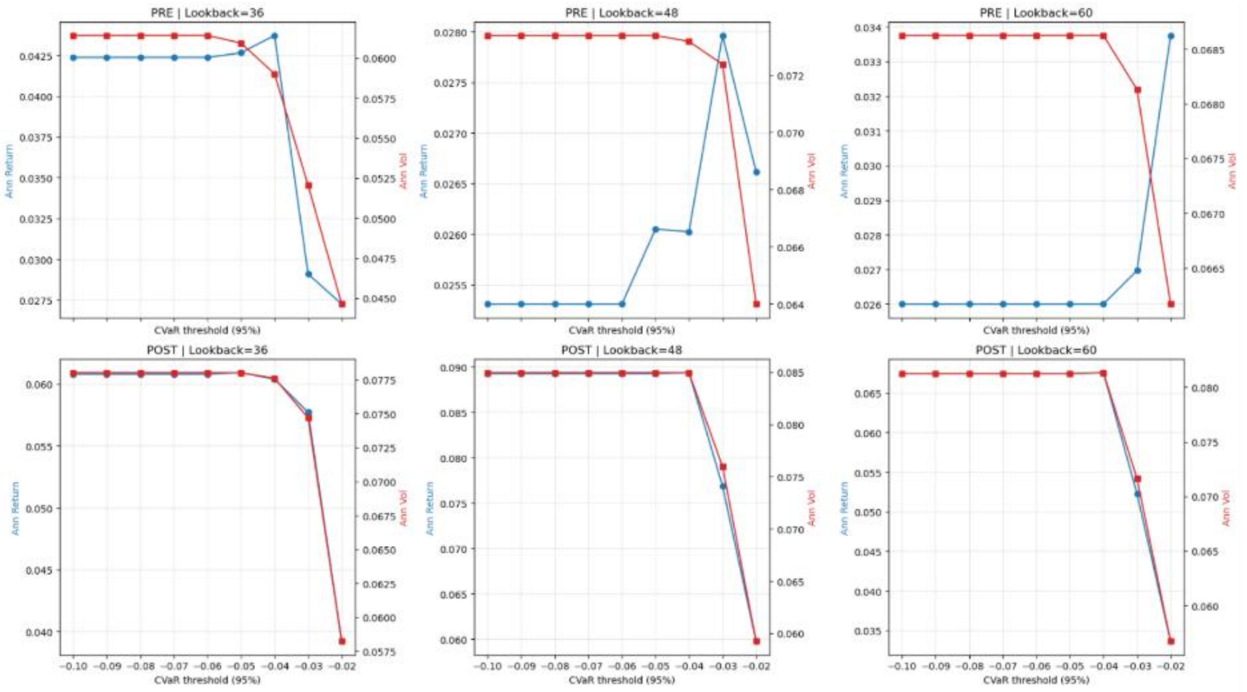

Six panels compare optimized portfolio performance across CVaR threshold levels (-10% to -2%) for Pre-Fukushima (top row, 2004-Feb 2011) and Post-Fukushima (bottom row, Mar 2011-2023) periods. Each row shows three lookback windows (36m, 48m, 60m). Blue lines with circular markers (left y-axis): Annualized return. Red lines with square markers (right y-axis): Annualized volatiliy. Top row (PRE): Returns remain relatively flat until stricter thresholds, with optimal Sharpe at -4% for 36m (0.74) but at -2% for 48m and 60m (0.42, 0.51). Volatility decreases as thresholds tighten. Bottom row (POST): Returns plateau from -10% to -4% at higher levels (36m: ~6.1%, 48m: ~8.9%, 60m: ~6.8%), then drop sharply. Optimal Sharpe ratios occur at moderate thresholds: -5% for 36m (0.78) and -4% for 48m/60m (1.05, 0.83). The shift in optimal CVaR thresholds from stricter levels (PRE) to more moderate levels (POST) suggests that overly restrictive downside constraints become counterproductive in the post-crisis environment, likely eliminating beneficial options strategies that help navigate increased volatility. Parameters: monthly rebalancing, Theoretical NEAR preference = 0.5, CVaR α = 95%, options cap = 30%, per-asset cap = 30%.

Figure 9. Optimized portfolio performance across different levels of near (risk/return preference)

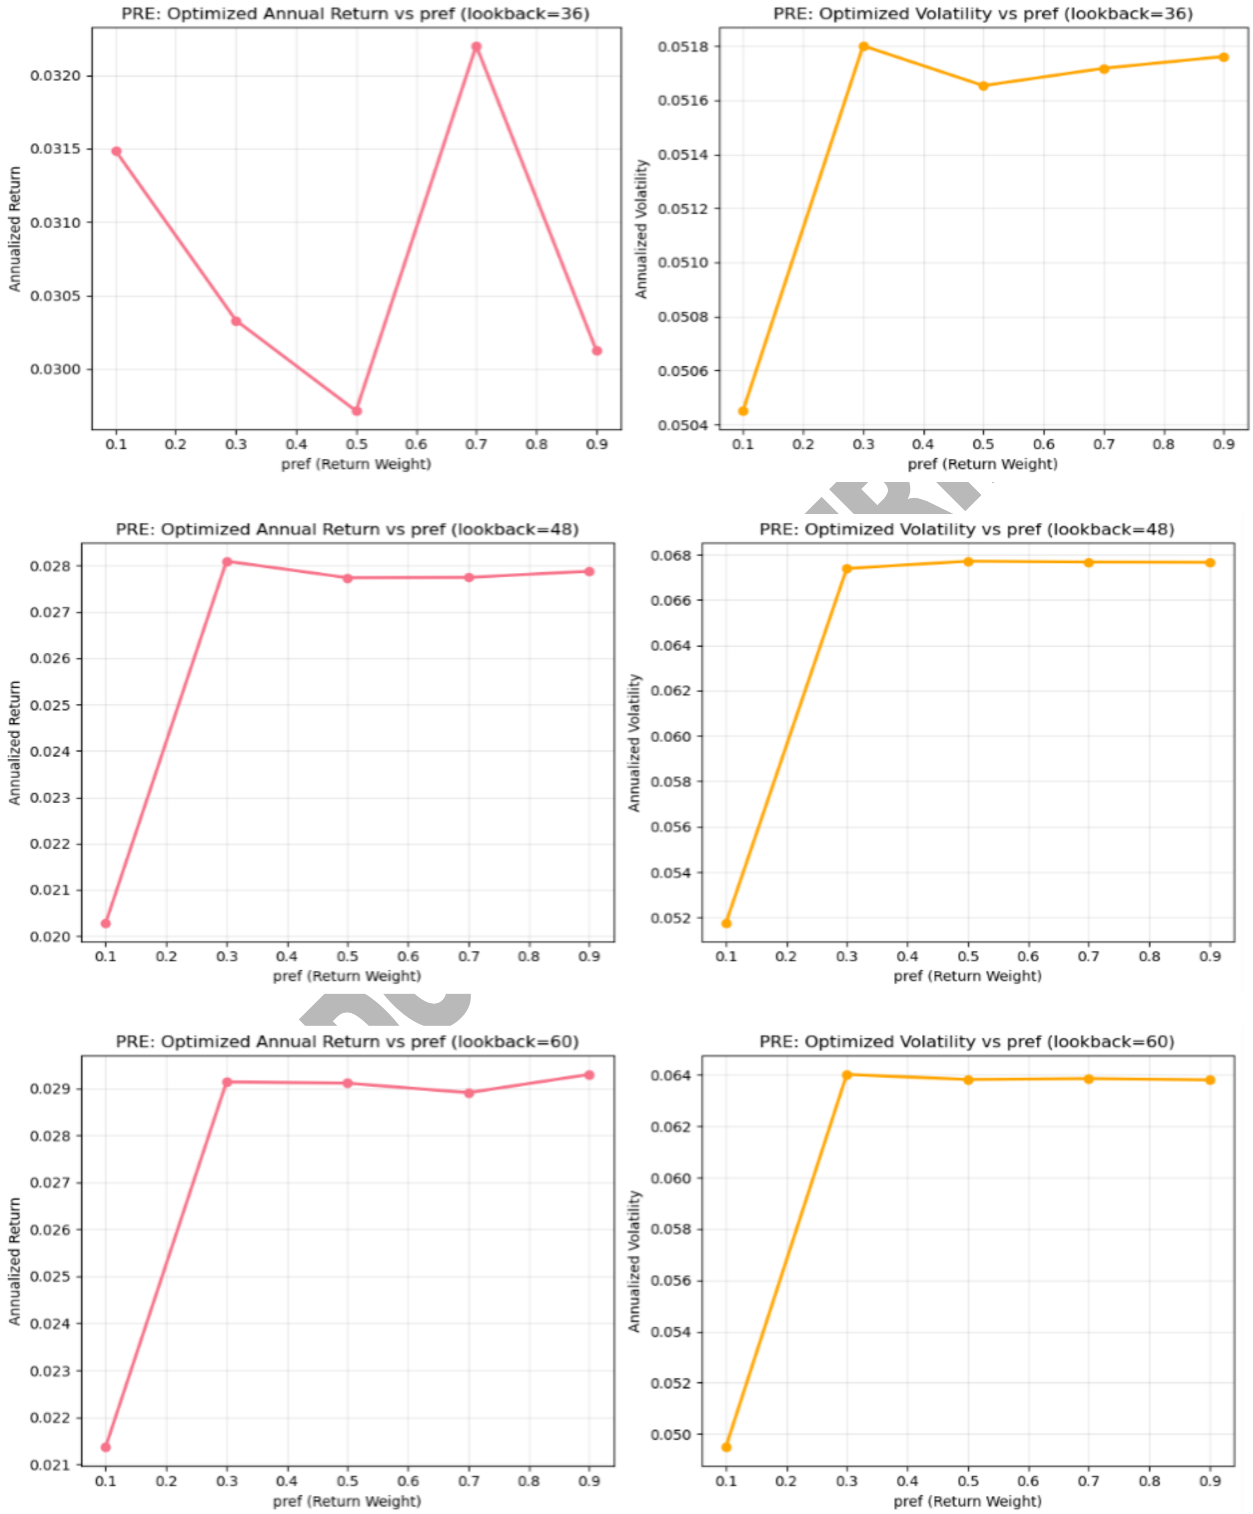

PRE: Optimized Annual Return vs pref (lookback=36)

PRE: Optimized Volatility vs pref (lookback=36)

PRE: Optimized Annual Return vs pref (lookback=48)

PRE: Optimized Volatility vs pref (lookback=48)

PRE: Optimized Annual Return vs pref (lookback=60)

PRE: Optimized Volatility vs pref (lookback=60)

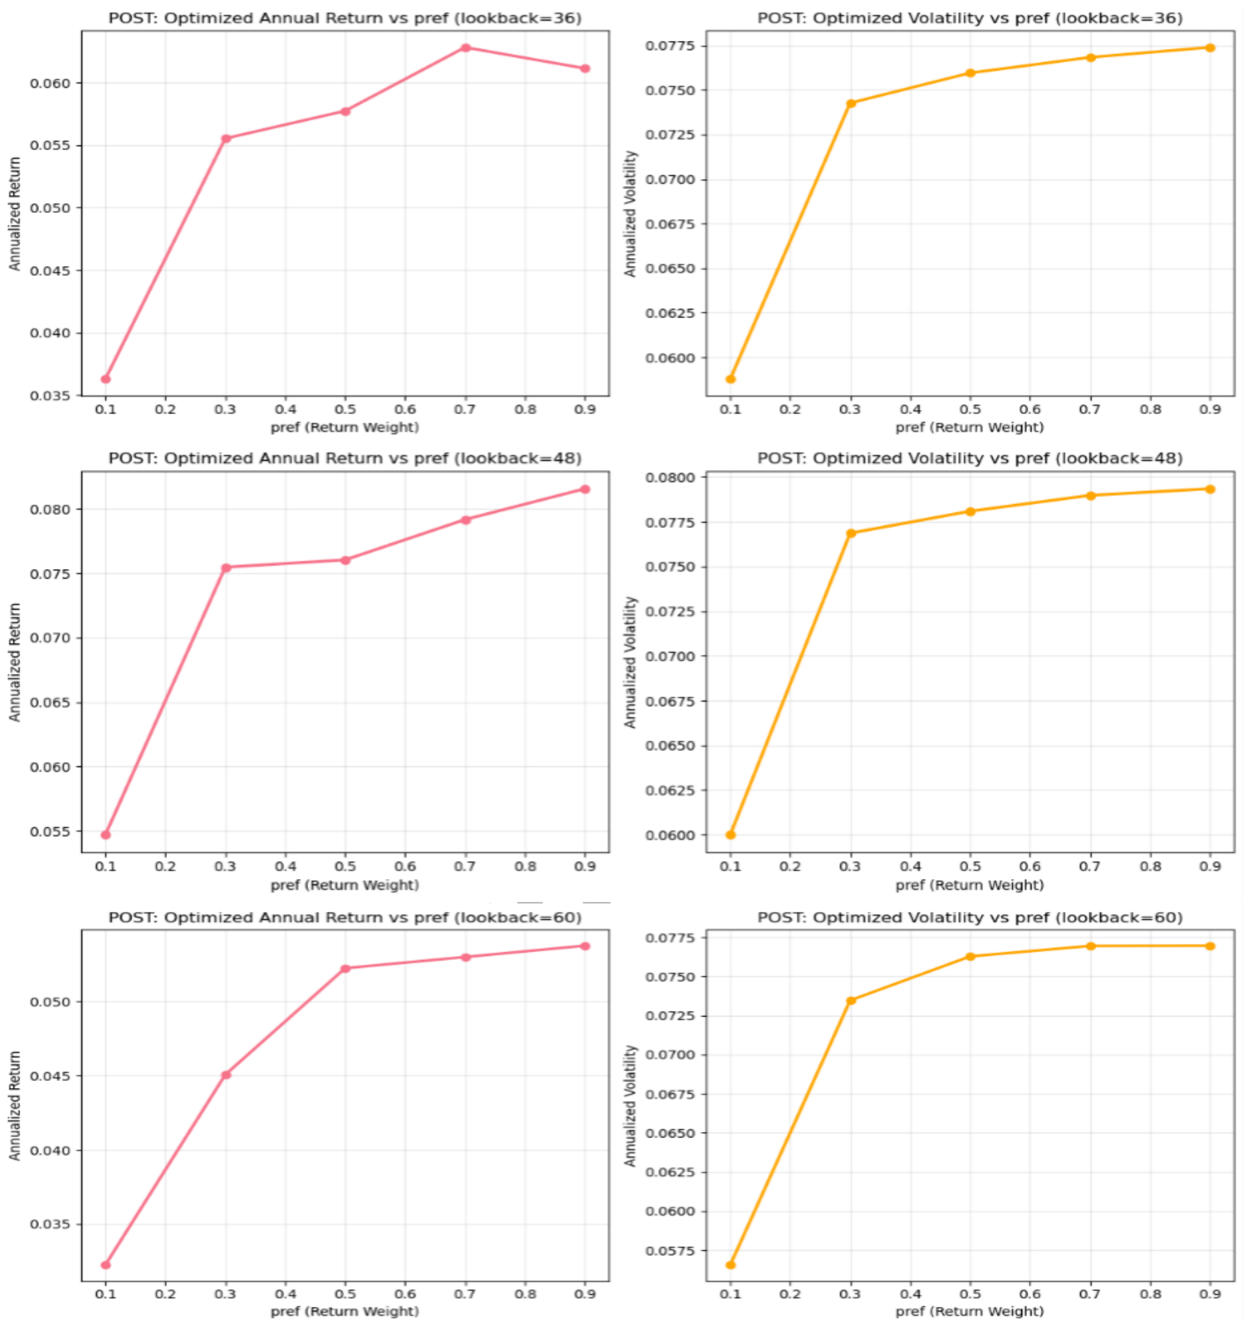

POST: Optimized Annual Return vs pref (lookback=36)

POST: Optimized Volatility vs pref (lookback=36)

POST: Optimized Annual Return vs pref (lookback=48)

POST: Optimized Volatility vs pref (lookback=48)

POST: Optimized Annual Return vs pref (lookback=60)

POST: Optimized Volatility vs pref (lookback=60)

Six paired panels compare optimized portfolio performance across return weight preference levels (0.1 to 0.9) for Pre-Fukushima (top three rows, 2004-Feb 2011) and Post-Fukushima (bottom three rows, Mar 2011-2023) periods. Each row shows a different lookback window (36m, 48m, 60m). Left panels: Annualized return (pink). Right panels: Annualized volatility (orange). Top rows (PRE): 36m shows V-shaped return pattern, dipping to ~2.97% at pref=0.5, recovering to ~3.22% at pref=0.7. Both 48m and 60m exhibit sharp return jumps at pref=0.3 (from ~2.0% to ~2.8-2.9%), then plateau. Volatility shows modest increases across all lookbacks, with steepest rise in 36m (5.05% to 5.18%). Bottom rows (POST): 36m shows steady return climb from ~3.6% to peak ~6.3% at pref=0.7. Both 48m and 60m display sharp return increases at pref=0.3 (48m: ~5.5% to ~7.5%, 60m: ~3.2% to ~5.3%), continuing upward thereafter. Volatility jumps significantly at pref=0.3 for all lookbacks (from ~5.7-6.0% to ~7.3-7.7%), then rises gradually. The consistent return jumps at pref=0.3 across longer lookbacks (48m, 60m) in both periods suggest this moderate return-weight level unlocks beneficial portfolio strategies. POST period shows higher absolute returns and steeper response to increased return preference, indicating greater opportunity for return enhancement in the post-crisis environment. Parameters: monthly rebalancing, CVaR α = 95%, options cap = 30%, per-asset cap = 30%.

Table 1: Summary Statistics for Stocks and Options

Panel A: Stocks

Variable	N	Mean	SD	Min	Max
Stocks (Daily)	43,945	0.001	0.040	-0.5	0.5
Stocks (Month-end)	2,160	0.003	0.046	-0.5	0.5

Panel B: Options

Calls (Daily)

Variable	N	Mean	SD	Min	Max
Return	120,909	-0.001	0.018	-0.5	0.5
Time-to-expiration	120,909	24.614	5.937	15	35
Moneyness	120,909	1.083	0.090	1.003	2.995
Delta	120,909	0.265	0.104	0.052	0.450

Puts (Daily)

Variable	N	Mean	SD	Min	Max
Return	157,370	0.001	0.014	-0.5	0.5
Time-to-expiration	157,370	24.602	5.959	15	35
Moneyness	157,370	0.929	0.058	0.491	1.000
Delta	157,370	-0.235	0.109	-0.450	-0.050

Calls (Month-end)

Variable	N	Mean	SD	Min	Max
Return	2,169	0.000	0.024	-0.243	0.5
Time-to-expiration	2,169	19.720	3.320	15	35
Moneyness	2,169	1.042	0.051	1.005	1.672
Delta	2,169	0.345	0.077	0.101	0.450

Puts (Month-end)

Variable	N	Mean	SD	Min	Max
Return	4,207	0.001	0.013	-0.2	0.304
Time-to-expiration	4,207	19.811	3.539	15	35
Moneyness	4,207	0.967	0.028	0.723	0.999
Delta	4,207	-0.317	0.094	-0.450	-0.069

Table 2: Initial Performance Comparison between Optimized and Equal-Weight Portfolios

	Optimized Portfolio	**Equal-weight Portfolio**
Total Return	129.90%	91.10%
Annualized Return	5.26%	4.07%
Annualized Volatility	8.33%	7.25%
Sharpe Ration	0.63	0.56
# of Obs (months)	195	195

This table provides initial assessment of the difference in performance between the optimized portfolio and the equal-weight equity-only benchmark portfolio

Table 3. Performance Comparison between Optimized and Equal-Weight Portfolios

Metric	Optimized Portfolio	Equal-Weight Benchmark
Tot. Return	129.90%	91.10%
Ann. Return	5.26%	4.07%
Ann. Vol.	8.33%	7.25%
Sharpe Ratio	0.631	0.561
Max. DD	-19.17%	-12.66%
Pos. Months	54.87%	55.90%
Obs. (months)	195	195

This table documents the performance comparison between the optimized portfolio (combining equities and OTM options with delta 0.1-0.4 using mean-variance optimization) and the equal-weight equity-only benchmark over the full sample period. The dependent variable is monthly portfolio return, derived from a 36-month rolling estimation window with monthly rebalancing. The optimized portfolio employs Ledoit-Wolf shrinkage for covariance estimation, maintains a 30% per-asset cap, limits options exposure to 30% of portfolio value, and enforces a CVaR (95%) constraint of -5%.

Table 4. Rebalancing Frequency Sensitivity

Freq.	**Opt SR.**	**Opt AR.**	**Opt Vol.**	**EW SR.**	**EW AR.**	**EW Vol.**	**Obs.**
1D	0.030	0.82%	27.37%	0.122	4.53%	37.05%	4,109
1W	0.313	4.68%	14.96%	0.404	6.09%	15.08%	860
1M	0.631	5.26%	8.33%	0.561	4.07%	7.25%	195
3M	-0.012	-0.05%	4.31%	0.084	0.36%	4.29%	65
12M	-0.293	-0.52%	1.77%	-0.084	-0.17%	1.96%	15

This table reports optimized portfolio performance across rebalancing frequencies using a 36-month lookback window. Theoretical NEAR = 0.5, options cap = 30%, and CVaR constraint = -5%. Frequencies range from 1D to 12M. SR = Sharpe Ratio; AR = Annualized Return; Vol. = Annualized Volatility; Opt. = optimized portfolio; EW = equal-weight benchmark; Obs. = number of rebalancing observations.

Table 6. Theoretical NEAR Parameter Sensitivity

Theo. NEAR	**Opt SR.**	**Opt AR.**	**Opt Vol.**	**Max. DD.**
0.0	0.403	1.43%	3.54%	-8.12%
0.1	0.646	4.09%	6.33%	-15.23%
0.2	0.659	4.85%	7.37%	-17.88%
0.3	0.637	4.93%	7.73%	-18.45%
0.4	0.622	5.07%	8.15%	-18.92%
0.5	0.631	5.26%	8.33%	-19.17%
0.6	0.633	5.31%	8.39%	-19.28%
0.7	0.633	5.35%	8.46%	-19.35%
0.8	0.640	5.41%	8.45%	-19.38%
0.9	0.640	5.42%	8.46%	-19.39%
1.0	0.637	5.38%	8.46%	-19.40%

This table presents portfolio performance across Theoretical NEAR parameter values using 36-month lookback, monthly rebalancing, options cap=30%, and CVaR=-5%. The Theoretical NEAR parameter α weights expected return vs variance in the objective function: maximize $\alpha \cdot E[R] - (1-\alpha) \cdot Var[R]$. When α=0, the optimizer focuses purely on variance minimization; when α=1, it focuses purely on return maximization. The dependent variable is Sharpe ratio.

Table 7. Options Cap Sensitivity

Op. Cap	**Opt SR.**	**Opt AR.**	**Opt Vol.**	**Max. DD.**	**Avg OW**
0%	0.561	4.07%	7.25%	-12.66%	0.00%
10%	0.598	4.52%	7.56%	-14.23%	8.7%
20%	0.615	4.89%	7.95%	-16.88%	17.3%
30%	0.631	5.26%	8.33%	-19.17%	24.1%
40%	0.628	5.31%	8.46%	-20.45%	31.8%
50%	0.619	5.28%	8.53%	-21.92%	38.2%
60%	0.605	5.19%	8.58%	-23.15%	44.5%

This table documents portfolio performance across options cap levels using 36-month lookback, monthly rebalancing, Theoretical NEAR =0.5, and CVaR=-5%. The options cap constrains the maximum aggregate portfolio weight allocated to call and put options. The dependent variable is Sharpe ratio. Average options weight shows the realized mean allocation to derivatives over the sample period.

is much larger, accounting for about 9.73 million records compared with about 44,000 equity entries. This reflects the fact that many utilities and uranium exploration firms in this group are actively traded in derivatives markets. Importantly, the dataset achieves very high levels of completeness, consistently above 95 percent, with nearly perfect coverage in more recent years.

The temporal profile of the data shows clear structural shifts. Option activity grew from just over 130,000 contracts in 2000 to nearly one million in 2022. Three phases are evident: a period of gradual increase from 2000 to 2007; heightened volatility during the global financial crisis in 2008–2009; and rapid growth beginning in 2018, likely reflecting changing energy market conditions and institutional interest in hedging energy price risks.

Equity prices in the nuclear dataset are right-skewed. The median share price is about $8, with most firms trading under $25, but a long tail extends above $100. This mix reflects the dual presence of small-cap uranium exploration firms and larger, established utilities. Options data reveal an evenly balanced market between put and call contracts, suggesting both bullish and bearish sentiment. The median contract has a time to expiration of around 79 days, indicating that traders tend to focus on short- to medium-term horizons. Most trading is clustered around near-the-money contracts, which are the most liquid and useful for risk management.

5.2 Brown Energy Companies

The brown energy dataset is narrower in size but broader in company coverage. It contains about 1.73 million stock records for 643 distinct tickers, encompassing coal, petroleum, natural gas, refining, and high-emission utility firms. Unlike the nuclear dataset, this sample contains

equity data only, but the coverage is extremely consistent, with fewer than 0.05 percent of observations missing key variables.

Price levels in this dataset are highly dispersed. The median share price is approximately \$20, the mean is around \$29, and the range extends from near zero to more than \$2,700. The bulk of observations cluster at lower price levels, consistent with the prominence of small- and mid-cap firms in fossil fuel extraction and related industries. At the same time, several large and stable energy companies are also represented, contributing to the long tail of higher valuations. Temporal coverage has been steady, with roughly 70,000 observations per year and about 6,000–7,000 per month in more recent years. The most frequently observed firms include long-standing participants in global energy markets such as VIA and SQM, each with more than 8,000 records, along with other oil, gas, and utility firms with consistent representation.

Taken together, the two datasets provide both breadth and depth for analyzing the financial characteristics of the energy sector. The nuclear dataset emphasizes the role of derivatives and highlights how investors actively manage risk in a low-carbon energy segment with relatively concentrated equity activity. The brown energy dataset provides broad coverage of fossil fuel and high-emission firms, offering a reliable longitudinal record of equity market dynamics in traditional energy industries. Both datasets are of high integrity, cover more than two decades, and together create a robust platform for comparing how financial markets treat low-carbon versus high-emission energy sources.

6 Empirical Methodology.

To demonstrate the return-enhancement capabilities of our approach using derivatives, we implement a portfolio optimization framework that combines a rolling-window mean–variance