Optimal Sequencing of Projects with Uncertain Regulatory Costs

A Technical Report submitted to the Department of Systems Engineering

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

Malek Thabet

Spring, 2025

Technical Project Team Members

Christopher Woods

Mitch Mitchell

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Roman Krzysztofowicz, Department of Systems Engineering

Optimal Sequencing of Projects with Uncertain Regulatory Costs

Christopher Woods*, Malek Thabet*, and Mitch Mitchell*

*Department of Systems and Information Engineering, University of Virginia, Charlottesville, Virginia 22904
Email: {gqf8uu, bmy4ww, jbm8efn}@virginia.edu

***Abstract*—Businesses often face uncertain costs when pursuing compliance with external requirements, such as regulatory mandates or industry standards. These uncertainties complicate long-term planning and require decision-makers to evaluate a wide range of operational and economic-political factors. This research, sponsored by CapTech—a Richmond, Virginia-based consultancy—contributes toward a decision-support tool for problems of this nature. In particular, it provides methods for specifying solvable optimization models in contexts where uncertainty must be quantified using expert judgment, available data, or a combination of both. This contribution consists of three components: (i) a stochastic optimization model, (ii) a judgmental forecasting procedure, and (iii) a statistical cost analysis.**

A stochastic optimization model forms the core of the decision system, selecting and sequencing multi-year mitigation projects and mitigation instruments to reduce a compliance metric to a target threshold with sufficiently high probability and to minimize expected total cost. To operationalize the model for a specific application, the uncertainty about its random inputs—future mitigation costs and the compliance metric's value at the end of the planning horizon—must first be quantified. In the case study presented here, the compliance metric is emissions, mitigation projects correspond to engineering or business initiatives that reduce emissions, and mitigation instruments correspond to emission offsets that can be purchased to meet a compliance threshold.

Due to the absence of data, expert judgment is solely used to quantify uncertainty about future project costs and future emissions. For each, a judgmental forecasting procedure elicits expert-assessed quantiles guided by relevant operational and economic-political factors. For emission offset costs, limited historical data exists and was used to identify distinct volatility regimes and to model distributions of cost change. They provide central credible intervals to guide the spread of expert-assessed quantiles, while expert's beliefs forecast the future volatility regime and the cost change direction.

Together, the three components form a decision system for compliance planning under uncertainty, the core engine of CapTech's decision-support tool.

***Index Terms*—Uncertainty quantification, stochastic optimization, judgmental forecasting, statistical analysis.**

I. Introduction

Businesses face increasing uncertainty in regulatory compliance costs due to shifting governmental policies, market volatility, and evolving industry standards. A recent report from *The Wall Street Journal* highlights how regulatory scrutiny of energy contracts has led companies to favor mitigation efforts that avoid such uncertainty, particularly when planning new power generation projects for data center demand [1]. More broadly, firms must balance these regulatory considerations alongside operational factors such as project feasibility, resource constraints, and internal strategic priorities. These circumstances underscore the necessity of rigorous methods to quantify and manage uncertainty in compliance planning—especially when relevant data does not exist or does not fully characterize the uncertainty.

Motivated by these challenges, CapTech, a consultancy based in Richmond, Virginia, sponsored this research to support the development of a decision-support tool for compliance planning under uncertainty. The tool is designed to assist businesses in selecting and sequencing multi-year projects to achieve a specified compliance target while minimizing expected total cost.

A particular decision system for compliance planning under uncertainty is developed, centered around a stochastic optimization model that is supported by methods for specifying random inputs. The model selects and sequences multi-year mitigation projects and allocates mitigation instruments over a planning horizon to achieve a specified compliance target with sufficiently high probability and minimum expected total cost. To apply the model to a specific problem, uncertainty about its random inputs—future mitigation costs and the value of the compliance metric at the end of the planning horizon—must first be quantified. This paper contributes two methods for this purpose: a judgmental forecasting procedure, used when data is unavailable or insufficient, and a statistical cost analysis, used when historical data provides useful but incomplete information. In the case study presented here, the compliance metric is emissions, mitigation projects correspond to engineering or business initiatives that reduce emissions, and mitigation instruments correspond to emission offsets that can be purchased to meet a compliance threshold.

The uncertainty about future project costs, future emission reductions, and future emission offset costs is quantified using expert judgment. For each random variable, experts assess three p-probability quantiles ($p = 0.1, 0.5, 0.9$) corresponding to the left tail, median, and right tail of its distribution. Variable-specific guidelines are provided to ensure that relevant economic-political and operational factors are taken into account during the assessment process. The assessed quantiles are used to fit candidate parametric distributions from three distribution families—Gumbel, logistic, and reflected Gumbel—selected to accommodate scenarios with right skew, symmetry, and left skew, respectively. The best-fitting distribution is then used to fully quantify the uncertainty about each random variable for input to the stochastic optimization model.

For future emission offset costs, limited historical data is available for Chicago Mercantile Exchange (CME) Global Emissions Offset (GEO) futures, which serve as a proxy for the emission offset spot market. While insufficient to forecast long-term price trends, this data was found to characterize market volatility across three distinct regimes with high, medium, and low volatility. Within each regime, the standardized daily rate of change forms an approximately stationary series, allowing parametric distributions to be fitted and central credible intervals to be calculated. These intervals provide users with historically consistent guidance for assessing the spread of expert-assessed quantiles for emission offset costs, while retaining full flexibility to specify the median based on internal beliefs about future market behavior.

The resulting stochastic optimization model is mathematically rigorous while maintaining accessibility for non-technical users, as it requires only inputs that can be specified through guided expert assessment and available data. It constitutes the core engine of a decision-support tool sponsored by CapTech, providing organizations with a principled approach to strategic planning in regulatory environments where uncertain costs complicate decision-making.

II. Stochastic Optimization Model

The core of the decision system is a stochastic optimization model designed to guide businesses in selecting and sequencing multi-year mitigation projects and allocating mitigation instruments under uncertain regulatory compliance costs. The model minimizes the expected total cost required to achieve a specified compliance target by the end of a planning horizon, while ensuring that the target is met with sufficiently high probability. To accomplish this, the model incorporates uncertainty about future mitigation costs and the future value of the compliance metric, which must be quantified prior to optimization. These uncertainties reflect the growth of project costs over time, deviations in the compliance metric from its historical trend, and changes in the cost of mitigation instruments. The model integrates these elements within a mixed-integer linear programming formulation, enabling practical solution with standard optimization software.

A. Inputs

The model requires several deterministic inputs related to the compliance metric, available projects, and mitigation instruments. The compliance metric value from the previous year is denoted by ϵ, and its historical mean annual growth rate is denoted by μ. The number of available projects is N_p, and project i has a duration of η_i years, a current total cost of w_i, and a total compliance metric reduction of γ_i. Project costs are incurred at the beginning of each year during execution, while compliance reductions are applied at the end of each year. The current cost to adjust reported emissions using mitigation instruments is w_c per unit.

Additional parameters define the planning horizon, compliance target, and constraints for a specific application. The planning horizon length is N_y years, and the target value for the compliance metric to be reached by the end of the final year is τ. The minimum acceptable probability of reaching this target is α. The maximum proportion of the expected total cost that can be spent in a single year is β. For operational or regulatory reasons, a minimum number of mitigation instruments κ_j may be required in year j, and a set of projects $\mathcal{P}_j$ may be manually specified to start in year j.

Uncertainty enters the decision system through three random variables that capture future deviations from current conditions. The annual growth rate of project costs is denoted by X_p, which multiplies project costs by a factor of $(1 + X_p)$ each year. The compliance metric grows multiplicatively each year by a factor of $(1+\mu+X_d)$, where X_d reflects deviations from the historical mean growth rate. This formulation captures an anchor-and-adjust heuristic, in which experts use the historical trend as a reference point when assessing uncertainty. The annual growth rate of the cost of mitigation instruments is denoted by X_c, which multiplies instrument costs by a factor of $(1 + X_c)$ each year.

B. Pre-Processing

Before solving the optimization problem, additional variables are derived to represent the behavior of projects, the compliance metric, and mitigation instruments within the finite planning horizon.

First, project durations and compliance reductions are adjusted to reflect any truncation that occurs when a project extends beyond the final year. If project i starts in year j, its effective duration within the horizon is $\eta_{ij} = \min\{\eta_i, N_y - j + 1\}$. The total compliance metric reduction from project i when started in year j is then $\gamma_{ij} = \gamma_i \eta_{ij}/\eta_i$.

Next, future project costs, required compliance reductions, and mitigation instrument costs are represented as random variables that incorporate the effects of uncertainty. If project i starts in year j, its total cost is given by

$$W_{ij} = \frac{w_i}{\eta_i} \sum_{k=1}^{\eta_{ij}} (1 + X_p)^{k+j-2}. \tag{1}$$

The total compliance reductions required by the end of the planning horizon to meet the target τ is denoted by W_t, with distribution function G, which is strictly increasing and continuous over its support.

The cost of mitigation instruments in year j is given by

$$W_{cj} = w_c(1 + X_c)^{j-1}. \tag{2}$$

The decision system requires expected values of future costs and a critical quantile of the required compliance reductions. These are pre-computed as follows. The expected total cost of project i started in year j is

$$\lambda_{ij} = \mathbb{E}(W_{ij}) = \frac{w_i}{\eta_i} \sum_{k=1}^{\eta_{ij}} \sum_{l=0}^{k+j-2} \binom{k+j-2}{l} \mathbb{E}(X_p^l). \tag{3}$$

The expected cost of mitigation instruments in year j is

$$\lambda_{cj} = \mathbb{E}(W_{cj}) = w_c \sum_{l=0}^{j-1} \binom{j-1}{l} \mathbb{E}(X_c^l). \tag{4}$$

Finally, the α-probability quantile of the required compliance reductions is $w_{t\alpha} = G^{-1}(\alpha)$. This quantile ensures that the compliance target is reached with the specified probability.

C. Mathematical Formulation

The decision system selects and sequences projects and allocates mitigation instruments over a finite planning horizon to minimize the expected total cost of achieving the compliance target with sufficiently high probability.

The model introduces binary decision variables a_{ij} to indicate whether project i is started in year j, for $i = 1, \ldots, N_p$ and $j = 1, \ldots, N_y$. A project is active in a given year if it was started in a previous year and has not yet completed its execution. The model also includes continuous decision variables a_{cj} to represent the number of mitigation instruments to be purchased in year j, for $j = 1, \ldots, N_y$. The expected total cost of the selected projects and mitigation instruments is represented by the continuous variable z.

The objective function minimizes the expected total cost:

$$\min\ z = \sum_{j=1}^{N_y} \left(\sum_{i=1}^{N_p} a_{ij}\lambda_{ij} + a_{cj}\lambda_{cj} \right). \tag{5}$$

To ensure the compliance target is reached with the specified probability α, the total compliance reductions from selected projects and mitigation instruments must meet or exceed the critical quantile $w_{t\alpha}$ of the required reductions random variable W_t:

$$\sum_{j=1}^{N_y} \left(\sum_{i=1}^{N_p} a_{ij}\gamma_{ij} + a_{cj} \right) \geq w_{t\alpha}. \tag{6}$$

To enforce budget discipline, the model limits the expected expenditures in each year to a fraction β of the expected total cost:

$$\sum_{i=1}^{N_p} a_{ij}\lambda_{ij} + a_{cj}\lambda_{cj} \leq \beta z, \qquad j = 1, \ldots, N_y. \tag{7}$$

Additional constraints allow the user to manually specify operational requirements. For example, a minimum number of mitigation instruments κ_j may be required in year j, enforced by:

$$a_{cj} \geq \kappa_j, \qquad j = 1, \ldots, N_y. \tag{8}$$

Similarly, a set of projects $\mathcal{P}_j$ may be specified to start in year j, enforced by:

$$a_{ij} = 1, \qquad \forall i \in \mathcal{P}_j,\ j = 1, \ldots, N_y. \tag{9}$$

Each project can only be started once over the planning horizon:

$$\sum_{j=1}^{N_y} a_{ij} \leq 1, \qquad i = 1, \ldots, N_p. \tag{10}$$

Finally, the binary nature of project selection is enforced:

$$a_{ij} \in \{0, 1\}, \qquad i = 1, \ldots, N_p,\ j = 1, \ldots, N_y. \tag{11}$$

This mathematical formulation provides a structured and rigorous approach for project sequencing and mitigation instrument allocation under uncertainty, while maintaining flexibility for user-specified constraints and operational requirements. The mixed-integer linear programming formulation ensures compatibility with standard optimization software and practical solution for real-world applications.

The case study presented henceforth considers emissions as the compliance metric of interest. Mitigation projects correspond to engineering or business initiatives that reduce emissions, while mitigation instruments correspond to carbon offsets that adjust reported emissions for compliance purposes.

III. Judgmental Forecasting Procedure

Experts will consider qualitative information, such as operational and economic-political factors, to quantify the uncertainty about X_p, X_d, and X_c. These inputs will be assessed using judgment provided in the form of three p-probability quantiles ($p = 0.1, 0.5, 0.9$), representative of the left tail, median, and right tail, respectively. Then, the quantiles will be used to fit models from the Gumbel, logistic, and reflected Gumbel families.

A. Guidelines and Instructions

Each of the three inputs will depend on the expert's ability to interpret various factors from both the internal and external environment. To support this process, users will first receive tailored guidelines for each input, providing key considerations specific to that input. In some cases, they may be asked to draw upon recent inflation trends and anticipated government policies, while in others they are asked to consider internal shifts in organizational development, leadership priorities or realignment in company strategy.

These tailored guidelines will serve as a foundation for users to make well-informed assessments. The following instructions will provide a step-by-step guide for assessing each quantile for each random variable, using the assessment of the annual growth rate of project costs X_p as an example.

- *0.5-Probability Quantile*: This represents the median estimate of the annual growth rate of project costs over the planning period. You may set a positive value if you believe the costs will increase, a zero value if you believe they will remain the same, and a negative value if you believe the costs will decrease. This quantile represents your best guess of a year-over-year project cost growth rate that is exceeded 50% of the time, or for 5 out of 10 years.
- *0.1-Probability Quantile*: Consider this an optimistic scenario where the annual growth rate of project costs will be less than anticipated. This quantile represents a year-over-year project cost growth rate that is exceeded 90% of the time, or for 9 out of 10 years.

- *0.9-Probability Quantile*: Consider this a pessimistic scenario where the annual growth rate of project costs will be greater than anticipated. This quantile represents a year-over-year project cost growth rate that is exceeded 10% of the time, or for 1 out of 10 years.

B. Parametric Distributions

The three quantiles provided by the expert only offer a crude pointwise representation of a continuous distribution function. However, these assessed quantiles suffice to estimate a parametric model which characterizes uncertainty about each input completely and conveniently. For each random variable, the judgmental parametric distribution functions can be estimated using the pointwise representation and the appropriate sample space [2].

Any distribution on the sample space which is the unbounded interval is an appropriate hypothesis for all three inputs. Furthermore, because assessed quantiles may not necessarily be centered around the median, three distribution families may be applied: Gumbel, logistic, and reflected Gumbel. The distribution function H for all three distributions is stated below.

1) *Gumbel Distribution*: Applicable when assessed quantiles are skewed to the right of the median,

$$H(x) = e^{-e^{-\frac{x-\beta}{\alpha}}}, \qquad -\infty < x < \infty. \tag{12}$$

2) *Logistic Distribution*: Applicable when assessed quantiles are roughly symmetric,

$$H(x) = \left(1 + e^{-\frac{x-\beta}{\alpha}}\right)^{-1}, \qquad -\infty < x < \infty. \tag{13}$$

3) *Reflected Gumbel Distribution*: Applicable when assessed quantiles are skewed to the left of the median,

$$H(x) = 1 - e^{-e^{\frac{x-\beta}{\alpha}}}, \qquad -\infty < x < \infty. \tag{14}$$

The quantile function of each distribution family—derived from the distribution function H—is used to obtain the location parameter β and the scale parameter α through the quantiles-moments (QM) method [2].

The goodness of fit of the model can be evaluated using the maximum absolute difference (MAD) between the pointwise representation and the hypothesized parametric distribution function. The hypothesized distribution with the best fit—the smallest MAD—will be selected. It is important to note that the MAD goodness-of-fit measure does not make any assumptions about the underlying distribution of the quantiles.

C. Demonstration

For illustrative purposes, consider a scenario in which an expert anticipates that the annual growth rate of project costs, X_p, will rise over the planning period due to recently implemented government policies and prevailing global developments. As a result, the expert may be inclined to assess the following quantiles.

1) *0.5-Probability Quantile*: 0.02.
2) *0.1-Probability Quantile*: 0.00.
3) *0.9-Probability Quantile*: 0.07.

The three distribution families are now fit to the quantiles using the QM method. Given the right skew nature of the assessment, the Gumbel distribution has the best fit. Fig. 1 shows the pointwise representation overlaid on the selected parametric distribution function.

D. Method for Computing Expectations

Once the parametric distribution functions for X_p and X_c have been selected, the expected values $\mathbb{E}(X_p^n)$ and $\mathbb{E}(X_c^n)$ must be calculated for $n = 1, \ldots, N_y$ so that λ_{ij} and λ_{cj} can be calculated for $i = 1, \ldots, N_p$ and $j = 1, \ldots, N_y$. These expected values are generally computationally expensive to calculate for large n, but they can be obtained with relative ease for the standard Gumbel and standard logistic distributions.

For any Gumbel or logistic variate X, the standard variate Z has parameters $(\beta, \alpha) = (0, 1)$ and can be obtained through the transformation $Z = (X - \beta)/\alpha$. Thus, the expected value $\mathbb{E}(X^n)$ for $n = 1, \ldots, N_y$ can be written in terms of $\mathbb{E}(Z^k)$ for $m = 1, \ldots, n$:

$$\mathbb{E}(X^n) = \sum_{k=0}^{n} \binom{n}{k} \beta^{n-k} \alpha^k \mathbb{E}(Z^k). \tag{15}$$

The standard Gumbel variate Z can be obtained from any reflected Gumbel variate X through the transformation $(\beta - X)/\alpha$, so the expected value $\mathbb{E}(X^n)$ for $n = 1, \ldots, N_y$ can be written in terms of $\mathbb{E}(Z^k)$ for $m = 1, \ldots, n$:

$$\mathbb{E}(X^n) = \sum_{k=0}^{n} \binom{n}{k} \beta^{n-k} (-1)^k \alpha^k \mathbb{E}(Z^k). \tag{16}$$

It is necessary to consider all information when quantifying uncertainty about these three random variables. In the case of X_c, historical market data is available and accessible, which can offer empirical guidance to assist the expert in their judgment.

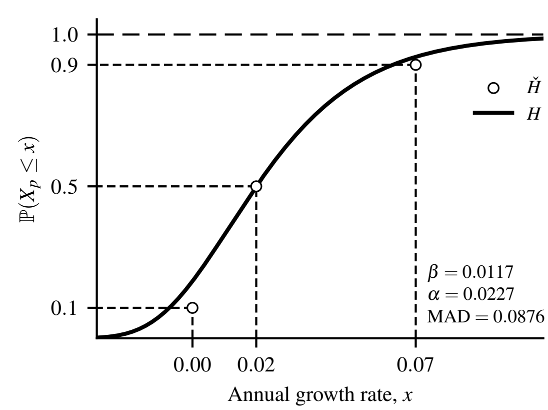

Fig. 1: The selected Gumbel distribution function H and pointwise representation $\check{H}$ of X_p from the demonstration.

IV. Statistical Analysis of GEO Futures

When a company wants to meet carbon footprint expectations or regulations, they can buy a Global Emissions Offset (GEO) futures contract from the Chicago Mercantile Exchange (CME) Group. The contract size for GEO futures is 1,000 offsets, with each offset representing the reduction of one metric ton of carbon dioxide or its equivalent in other greenhouse gases. GEO futures contracts are listed monthly for the current year and at least the next three calendar years.

The dataset comprises front-month Global Emissions Offset Futures contracts from the Chicago Mercantile Exchange, spanning March 1, 2021, to January 28, 2025. It includes trade dates, daily opening, highest, lowest, and settled prices, as well as trading volume. Exploratory data analysis was conducted to visualize price trends and variability over time.

A. Exploratory Data Analysis

The exploratory data analysis involved graphing key time series aspects, including closing price, volume, daily price change (in dollars), and daily rate of change. This analysis revealed three separate regimes over its history, corresponding to high, medium, and low volatility. The data were broken into regimes according to the following date ranges.

- *High*: 7/6/2021–8/12/2022, $N = 280$ trading days.
- *Medium*: 8/15/2022–7/28/2023, $N = 240$ trading days.
- *Low*: 7/31/2023–4/25/2024, $N = 187$ trading days.

The tails of the data (3/1/2021–7/5/2021 and 4/26/2024–1/28/2025) were excluded from statistical analysis due to low volume. For each regime, the four time series (closing price, volume, daily price change, and daily rate of change) were visualized. Additionally, for each price change unit [$, rate] within each regime, the sample median $x_{0.5}$, sample mean m, and sample standard deviation s were calculated (TABLE I). The data were then standardized using $z(n) = (x(n) - m)/s$ for $n = 1, \ldots, N$, and the standardized series were plotted.

B. Fitting Distributions

In Fig. 2, the closing price series lacks stationarity, but when daily price rate of change is standardized using the methods listed above, the resulting series demonstrates stationarity. This stationarity allows the standardized daily price rate of change to be treated as a random variable and be modeled using parametric distributions.

The web-based distribution fitter, DFit, which accompanies the book by Krzysztofowicz [2], enables the user to fit parametric distributions to a random sample. The user uploads the sample, the bounds on the sample space (lower, upper, both, or none), and the program provides possible distributions that could fit the data. It then fits these distributions, estimates their parameters, calculates the maximum absolute difference (MAD) between the fitted parametric distribution and the empirical distribution, and performs the Kolmogorov–Smirnov (K–S) goodness-of-fit test. The user is then able to superimpose the fitted distributions onto a graph of the empirical distribution function in order to visually assess the goodness of fit.

TABLE I: The sample median $x_{0.5}$, sample mean m, and sample standard deviation s of the daily rate of change of the price of GEO futures for each of the volatility regimes.

Volatility Regime	**High**	**Medium**	**Low**
Daily Price Change [$]			
Median $x_{0.5}$	0.0	-0.0100	0.0
Mean m	0.0035	-0.0123	-0.0029
Standard deviation s	0.2132	0.1222	0.0623
Daily Rate of Change			
Median $x_{0.5}$	0.0	-0.0042	0.0
Mean m	0.0017	-0.0034	0.0004
Standard deviation s	0.0383	0.0668	0.0941

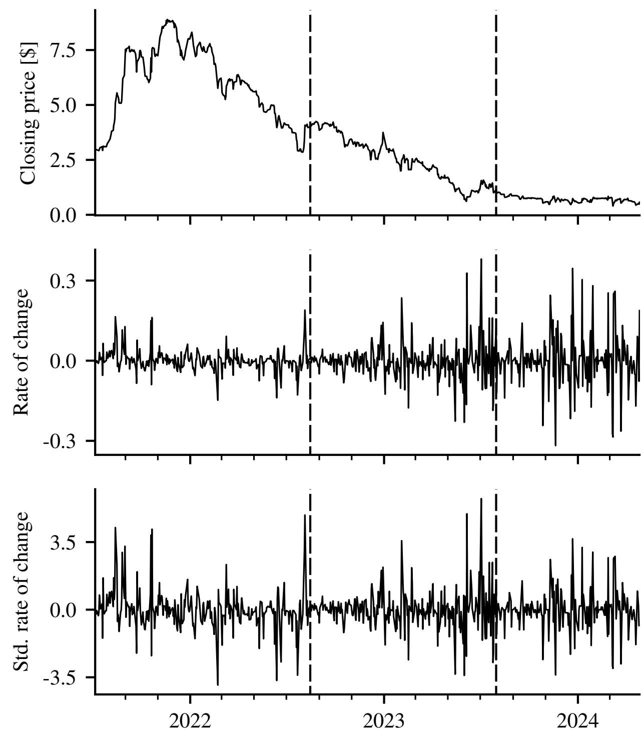

Fig. 2: The daily closing price, rate of change, and standardized rate of change of GEO futures during the three volatility regimes (high, medium, and low, from left to right).

Exploratory data analysis showed that the standardized daily rate of change and the standardized daily price change had similar variability across three volatility regimes.

Each of these six samples (three regimes for two metrics) was analyzed as well as one combined sample of all three regimes for each metric. Assuming an unbounded sample space, five distributions—logistic, Laplace, Gumbel, reflected Gumbel, and normal—were estimated and their goodness-of-fit was evaluated. (The equations for these distributions can be found in [2]).

Judging each distribution based on the MAD and K–S statistic, the Laplace and logistic fit all eight samples exceptionally well, better than the other three distributions. Laplace had the lowest MAD and K–S statistic for all eight samples.

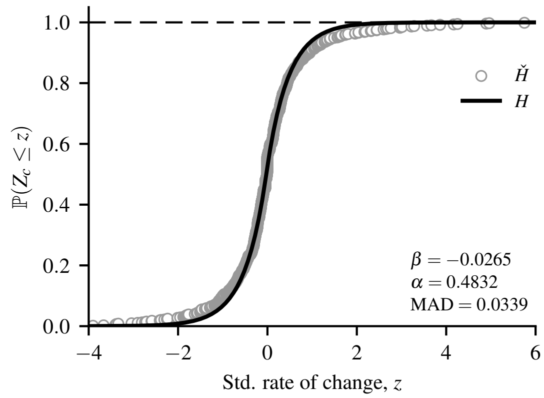

Fig. 3: The selected Laplace distribution function H and empirical distribution function $\breve{H}$ of the standardized daily rate of change Z_c for the combined sample of all volatility regimes.

Based on these results, the Laplace distribution was selected to characterize the variability of the standardized daily price change and the standardized daily rate of change (Fig. 3).

C. Calculating Central Credible Intervals

The statistical analysis of GEO futures prices is used to provide experts with a 0.8-probability central credible interval (CCI) for each price volatility regime. Experts forecast the volatility regime and use the corresponding CCI as an aid to assessing quantiles. This method adopts the expert-assessed median exactly without guidance, but informs the expert of the width of the 0.8-probability CCI based on historical data. This step should help to prevent a gross underestimation or overestimation of uncertainty in the forecasted volatility regime.

The previous analysis showed that the Laplace distribution fitted to the combined sample of the standardized daily rate of change should be used to provide this guidance, but it needs to be transformed to better match the variables introduced in Section II.

The standardized rate of change of carbon offset credit costs Z_c represents historical fluctuations in the front-month GEO futures contracts at the CME and is computed as $Z_c = (X_c - m)/s$, where m and s are the sample mean and standard deviation of X_c within a given volatility regime. The Laplace distribution function H of the standardized daily rate of change is shown in Fig. 3 and its expression is:

$$H(z) = \begin{cases} \frac{1}{2}\exp\left(\frac{z+0.0265}{0.4832}\right) & \text{if } z < -0.0265; \\ 1 - \frac{1}{2}\exp\left(-\frac{z+0.0265}{0.4832}\right) & \text{if } z \geq -0.0265. \end{cases} \tag{17}$$

The quantile function for the Laplace distribution was then used to calculate the prior 0.8-probability CCI as $z_{0.9} - z_{0.1} = 1.56$. The width of the 0.8-probability CCI for each regime was then calculated as $x_{0.9} - x_{0.1} = s(z_{0.9} - z_{0.1})$.

TABLE II shows that the 0.8-probability CCI for the low volatility regime is the widest, reflecting the highest degree of uncertainty. In contrast, the CCI for the high volatility regime is the narrowest, reflecting comparatively lower uncertainty. This is seemingly paradoxical, but is in fact easily explained. As shown in Fig. 2, the price is highest during the high volatility regime, followed by the medium volatility regime, and then the low volatility regime. As such, seemingly small fluctuations during the low volatility regime are actually larger relative changes than those experienced during the high volatility regime. Thus, it can reasonably be interpreted that high volatility corresponds to times of high economic-political activity towards emissions where absolute changes in price are large but relative changes are low, and vice versa for low volatility.

TABLE II: The sample standard deviation s of the daily rate of change of the price of GEO futures and the width of the prior 0.8-probability CCI $x_{0.9} - x_{0.1}$ for each of the volatility regimes.

Volatility Regime	**High**	**Medium**	**Low**
Standard deviation s	0.0383	0.0668	0.0941
$x_{0.9} - x_{0.1}$	0.0596	0.1039	0.1463

V. Summary and Conclusions

This paper presents a decision system that integrates judgmental forecasting and statistical cost analysis to quantify the uncertainty about the inputs to a stochastic optimization model to support compliance planning. The model fully quantifies the uncertainty about future project costs, compliance metric growth, and mitigation instrument costs to minimize the expected cost of achieving a compliance target.

Expert judgment is elicited through a structured forecasting procedure, in which users assess three quantiles for each uncertain input variable. These assessments are used to estimate parametric distributions from selected families, enabling complete quantification of uncertainty. Where available, historical data supplements the expert's judgment—statistical analysis of carbon offset futures provides credible intervals to guide quantile assessment for the cost of mitigation instruments.

Formulated as a mixed-integer linear program, the stochastic optimization model integrates these inputs and selects an optimal sequence of projects and mitigation instruments to achieve the compliance target with sufficiently high probability. This approach enables CapTech's decision-support tool to combine data-driven insights with expert knowledge, providing organizations with a principled framework for regulatory cost planning.

Acknowledgments

The authors would like to thank CapTech consultancy and their representatives Wray Mills, Whitney Olmsted, Molly Nesemann, Meagan Haupt, and Raymond Rosalez for their sponsoring and defining this project, and to Professor Roman Krzysztofowicz for providing technical guidance.

References

[1] J. Lee, "AI fever in power stocks moves from nuclear to plain natural gas," *The Wall Street Journal*, 1 March 2025.

[2] R. Krzysztofowicz, *Probabilistic Forecasts and Optimal Decisions*. Wiley, 2025.