Міністерство освіти і науки України

Національний університет «Києво-Могилянська академія»

Факультет економічних наук

Кафедра фінансів

Магістерська робота

ОСВІТНІЙ СТУПІНЬ - МАГІСТР

на тему: «**CONSTRUCTION OF AN INFORMATION-SHOCK-RESILIENT PORTFOLIO: INTEGRATION OF GARCH FORECASTS AND LLM SIGNALS WITHIN THE BLACK-LITTERMAN FRAMEWORK - ПОБУДОВА СТІЙКОГО ДО ІНФОРМАЦІЙНИХ ШОКІВ ПОРТФЕЛЯ: ІНТЕГРАЦІЯ GARCH ПРОГНОЗІВ ТА СИГНАЛІВ LLM У РАМКАХ МОДЕЛІ БЛЕКА-ЛІТТЕРМАНА** »

Виконав: студент 2-го року навчання,
Спеціальності:
072 «Фінанси, банківська справа, страхування та фондовий ринок»
Тарасов Арсен Андрійович

Керівник: Семко Р.Б.
кандидат економічних наук, доцент
Рецензент Григор’єв Г. С.
доктор економічних наук, доцент

Магістерська робота захищена
з оцінкою «_______________________»

Секретар ЕК ______________________
«____»_______________202__р.

Київ 2026

ЗМІСТ

INTRODUCTION

Relevance of the research topic. Given the global and highly interconnected economy of the 21st century, one could say that the financial markets are more volatile than ever and tend to evolve structurally. Such a statement is true and it has to do with a multitude of economic, geopolitical, and technological factors. Therefore, a market environment where traditional methods of portfolio management have been found to be inapt in solving the market-related issues. The crisis of 2021-2024 is just another example of the shortcomings of traditional approaches towards portfolio management, heavily reliant on the mean-variance paradigm, in being blind to tail risk, which is still very relevant to the current market conditions.

Moreover, in a world where information is king, the value of both traditional and non-traditional assets is very sensitive to news and prevailing market sentiment. Such a phenomenon may lead to an initial overreaction that eventually cools down. Therefore, today more than ever, there is a need for dynamic portfolio management strategies that are capable of quick adjustments depending on the situation. At the moment, there is a lot of attention being paid to portfolio management methods that allow for swift adjustment to the market conditions.

A step further from the mean-variance paradigm is the Black-Litterman model, which allows for improved volatility predictions through the use of GARCH models and incorporates more sophisticated sentiment metrics, such as our Shock Alpha metric, which measures market pessimism or euphoria.

After the crisis in Ukraine and thanks to what has happened in other emerging economies, the concept of antifragile investing is no longer a matter of theory and has become a must. The main goal of the research remains to further develop an effective algorithmic model.

The theoretical and methodological foundation of portfolio investment and advanced risk management has been the focus of major foreign researchers, such as H. Markowitz, F. Black, R. Litterman, W. Sharpe, and R. Engle, who introduced the GARCH methodology. Local researchers like A.B. Kaminskyi, L.O. Prymostka, V.V.

Vitlinskyi, and O.I. Baranovskyi have equally made important contributions to the development of financial analysis tools and volatility models with a special emphasis on emerging markets. On the contrary, adaptive models' usage, especially the Black-Litterman model, for extremely volatile and risky instruments like cryptocurrencies traded in an environment of informational and systemic shock extreme stress, is a field still under-researched. Addressing this scientific challenge forms the core of this article by delving into it extensively [1-4][26-\32].

Purpose and tasks of the research. The primary aim of the present research is to develop, validate and empirically examine the most efficient methodology for managing an asset portfolio. The principal way to do this is, on the one hand, to forecast volatility using GARCH modeling, and on the other hand, to incorporate proprietary "Shock Alpha" sentiment signals within the Black-Litterman framework, in order to have a combination that effectively compensates for and neutralizes the harmful effects of sudden market shocks. Besides, the paper intends to compare the newly created resilient portfolio and its risk-adjusted return profile with standard investment benchmarks, primarily the classical Markowitz portfolio.

In order to accomplish this far-reaching target, the following subordinate objectives have been defined and fulfilled in the context of the dissertation:

- To thoroughly study the theoretical basis of the Black-Litterman model and critically evaluate its structural merits over standard optimization methods;
- To draw up a sample of 50 representative assets from multiple classes (cryptocurrencies, traditional stocks, and commodities) and deeply analyze their price behavior for the turbulent 2021-2024 period;
- To conceptually and mathematically formulate a measure of ex-post asset resilience called "Shock Alpha" metric and then experimentally verify their relative strength in withstanding news shocks;
- To develop code for three different portfolio strategies: a classical approach (Markowitz MinVar), an agnostic approach (Black-Litterman with GARCH

forecasts but without sentiment signals), and a fully resilient strategy (incorporating LLM-simulated signals);

- To test the performance of the mentioned portfolios on the 2024 out-of-sample data by using a wide range of financial metrics including the Sharpe ratio, the Sortino ratio, and the Maximum Drawdown;
- To implement a comprehensive web-based visualization platform that supports full interactivity, using the R Shiny framework, primarily for the presentation of the dynamic model results and sophisticated scenario analysis.

Object and subject of the research. The object of the research is the broad process of managing a very diversified investment portfolio, which strategically blends the modern cryptos and the traditional financial assets. The subject of the research covers the particular mathematical techniques and algorithmic tools for portfolio optimization, which are primarily based on the Black-Litterman model and, at the same time, considerably advanced by the GARCH volatility modeling and the quantitative news analysis of news-driven shocks.

Research methods. To effectively carry out the tasks specified, this study utilizes a well-rounded mix of general scientific and specialized empirical methods. Among them:

- The systemic approach together with the theoretical generalization, helped to thoroughly analyze the current paradigm of portfolio investing and to recognize the limitations of classical theories;
- Advanced econometric methods (mainly referring to the GARCH class of models), were used to forecast asset volatility with high precision and to assess market risks on a highly dynamic basis;
- Bayesian optimization methods (which are the basis of the Black-Litterman model), were used to find the best combination of the portfolio components by smartly integrating the market forecasts oriented to the future;

- Detailed and sound statistical tests together with scenario models were rigorously applied to evaluate strictly the portfolios' capability of withstanding the most severe exogenous shocks, mainly by means of the "Shock Alpha" measure;
- Picture and chart methods, were used to present the complicated results of the empirical research in an extensive way inside the interactive R (Shiny) environment.

Scientific novelty of the obtained results. The essence of the novelty of the results obtained is the successful theoretical and practical development of a very complete, multi-tiered investment portfolio management approach that is specifically aimed at handling extreme market uncertainty effectively. The scientific novelty of the main points is defined as follows:

- This study is the first to recommend the direct use of a specially-designed resilience measure to news shocks (the "Shock Alpha" index) as one of the components of Investor Views in the Black-Litterman model framework. This new inclusion makes it possible to change asset weights very quickly when the market experiences a high level of turbulence;
- The return forecasting technique for future cryptocurrency as well as traditional assets has been notably advanced by the combination of GARCH volatility modeling and Bayesian optimization. The classic Markowitz model, on the other hand, does not have this advanced integration that considers the characteristic "heavy tails" of the empirical return distributions very effectively;
- The R programming language, in practice, has been taken to a farther level of sophistication with an interactive investment decision support system (the Shiny application) at its core. This system makes it possible to perform stress tests of complicated strategies instantly under the usage of the real historical data of the extremely difficult period of 2021-2024.

Practical significance of the obtained results. The practical side of the paper is absolutely clear from the fact that the authors managed to make it happen that were previously only conceptual phases. These phases can be briefly called the design of a resilient algorithmic core and its corresponding implementation in the R programming language. Using these kinds of "anti-fragile" portfolios, it is easy for individual retail investors as well as big institutional players to construct them effortlessly. In addition, the full-scale approach being introduced here is likely to be used by the advanced investment institutions and professional financial analysts as an effective tool to systematically reduce the losses in their portfolios under the unforeseen times of very extreme geopolitical and macroeconomic shocks.

CHAPTER 1.
THEORETICAL AND METHODOLOGICAL FOUNDATIONS OF PORTFOLIO MANAGEMENT UNDER MARKET SHOCKS

1.1 Evolution of portfolio theory: from Markowitz to Black-Litterman

Modern quantitative finance, including the scientific management of portfolios, is often traced back to the work of Harry Markowitz. His 1952 publication titled "Portfolio Selection" formally presented the Modern Portfolio Theory (MPT), marking the beginning of this era. Before the revolutionary discovery of Markowitz, investors mainly concentrated on fundamental analyses of single securities, strictly evaluating them individually to find intrinsically undervalued or overvalued assets. Markowitz did away with this disjointed pattern by providing a mathematical description of diversification and redirecting the study from isolated financial items to the portfolio of investments as a whole, an integrated entity [5].

The theoretical framework of current portfolio management has lately evolved towards the use of unstructured data, with Lopez-Lira and Tang (2023) setting a significant standard by proving that Large Language Models (LLMs) carry a great potential for predicting stock returns based on sentiment analysis. Further, Li and Tan (2022) confirmed that it is possible to use machine learning for creating "views" for the Black-Litterman model, thus they successfully connected the dots between algorithmic forecasting and portfolio building. Nevertheless, rather than viewing sentiment as a variable factor and ignoring the changes in market risk over time, especially at times of severe information shocks, these research works limit the usage of AI-based indicators to just one point in time.

Latest research in this area has introduced quite a few solid ideas. For example, the multi-LLM sentiment aggregation method of Mantshimuli and Mwamba (2025) is quite an interesting development, as it is designed to lessen the biases of individual models in Black-Litterman views. On a similar note, Lee et al. (2025) have looked into how using LLMs directly may improve the accuracy of day-to-day views inside the Bayesian framework. Besides that, Iacovides et al. (2024) brought out "FinLlama, " a financial sentiment analysis tool for an automated trading environment, which really helps in understanding the importance of sector-specific semantic knowledge. Still, the majority of these present-day systems still concentrate on finesse the "Views" part and tend to forget the quantitatively necessary "Volatility Radar" we have to use to be able to handle changing market conditions. Normally, these models are based on either historical measures of risk or are simple sentiment scores and in reality these two approaches do not reflect the "Volatility Clustering" that is typical of real-market shocks[16][18-20].

This paper sets itself apart from those researches mentioned above by designing a single "Resilient" system which combines qualitative semantic context with high-frequency econometric forecasting. Contrary to the models offered by Mantshimuli and Mwamba (2025) or Lee et al. (2025), this one incorporates a daily GARCH(1, 1) framework to deliver a continuously updated estimate of conditional variance. Besides, the "Shock Alpha Index" is an innovative feature of the paper that has been calibrated using 37 different historical crises illustrated in the Appendix A which facilitate the portfolio to get past simple resistance to real anti-fragility. By algorithmically adjusting the confidence matrix through this double-layer perspective, the suggested model not only safeguards capital during periods of downturn but also reallocates it to take advantage of market disorder, which is a missing ability in the literature that exclusively focuses on sentiment or static optimization.

At the heart of classical MPT lies the mean-variance optimization model. This model suggests that a rational, risk-averse investor is always looking to increase the expected return from the portfolio to the maximum for a certain amount of risk, or on the other hand, reduce risk to the minimum for the level of expected return set. In this classical sense, risk is being equated to the variance (or standard deviation) of the asset returns, while the complex relationship among different assets is expressed through their mathematical covariances. Thanks to such a detailed definition, it was possible to develop the concept of the "efficient frontier", a hypothetical range of optimal portfolios that deliver the highest expected return for any given level of risk. The classic depiction of this beautiful yet delicate concept is shown below.

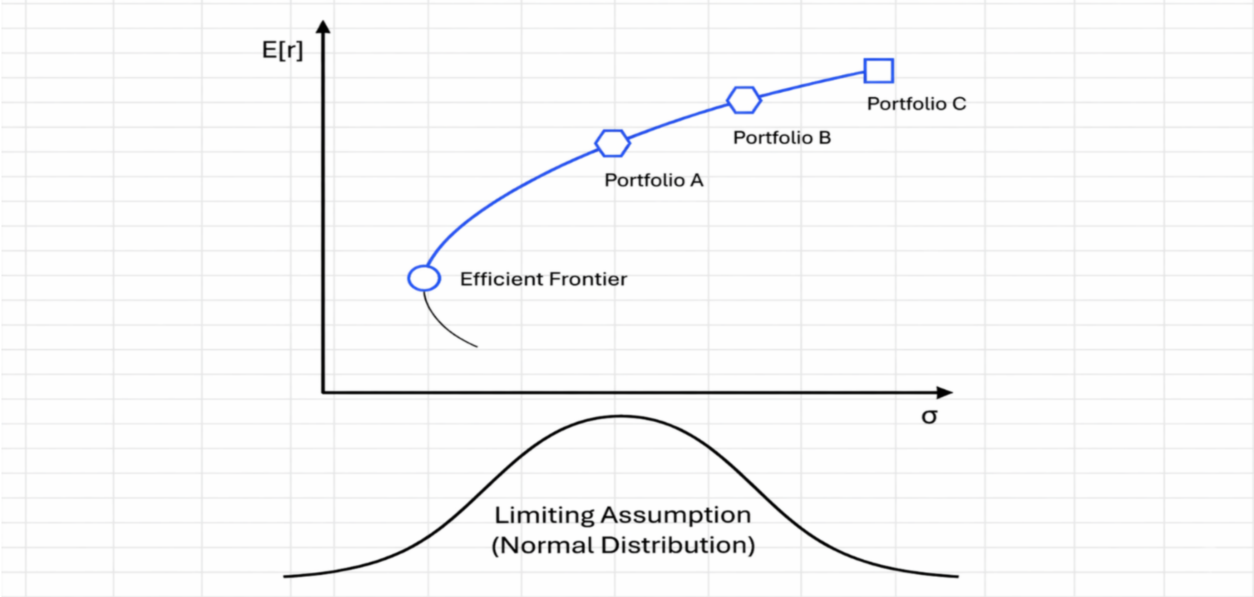

Fig. 1.1 - Visual Representation of Modern Portfolio Theory (MPT) and its Limiting Assumption.

This abstract visualization conveys the strict mathematical beauty of Markowitz's mean-variance optimization. The main “Efficient Frontier” is depicted as a delicate, see-through sculpture made of data. A line of perfectly arranged geometric data points (stands for portfolios A, B, and C) is accurately aligned to this curve by their Expected Return (E[r]) and Risk/Variance (σ). Most importantly, the basic limiting assumption, that returns are normally distributed, is illustrated by a flawless mathematical bell curve that is visibly cracked and breaking apart at the extremes. This fragmentation, on the one hand, shows the visual representation of the weakness of MPT when subjected to extreme market shocks, on the other hand, it systematically challenges the assumptions of MPT and makes the efficient frontier unstable[6].

Regardless of the profound theoretical elegance and the foundational role it plays in financial economics, the Markowitz mean-variance optimization model has been confronted with severe and well-documented limitations in its practical, real-world application. The highest criticism of MPT is the fact that it is very sensitive to its primary input parameters: expected asset returns and the covariance matrix. Because expected returns are exceptionally difficult to predict with any level of statistical accuracy, the classical Markowitz optimizer often becomes an "error maximizer." It fervently invests the major part of the capital in assets whose historical return estimates are overly optimistic and that are negatively correlated, which results in very concentrated, counter-intuitive, and practically unviable portfolios. Besides, the classical MPT relies heavily on the fragile statistical assumption that asset returns are normally (Gaussian) distributed. This assumption continually gives a lower estimate of the probability extreme market events, called by some "fat tails" or "black swan" events, thus the classical model is inherently flawed during global financial crises situations as happened in the severely geopolitical and macroeconomic shocks observed in the 2021-2024 period[7].

In order to systematically counteract the severe practical flaws the Markowitz framework has, Fischer Black and Robert Litterman, who were at Goldman Sachs in 1990, came up with the Black-Litterman (BL) model. The BL model is a giant evolutionary step in portfolio theory. It effectively changes the quantitative paradigm from pure historical data optimization to a very sophisticated Bayesian approach. Thus, they recognized that the past average returns are poor predictors of the future performance, and the Black-Litterman model smartly solves the major input problem of MPT. Instead of forcing the investor or analyst to supply absolute expected returns for all assets, an almost impossible task, the BL model starts off the optimization with a strong, highly stable, and neutral point: the implied equilibrium returns. These implied market returns are mathematically backed by the global market capitalization weights of the assets through the Capital Asset Pricing Model (CAPM), essentially assuming that the global market portfolio is already efficiently priced. The core innovation of the Bayesian approach is shown visually below.

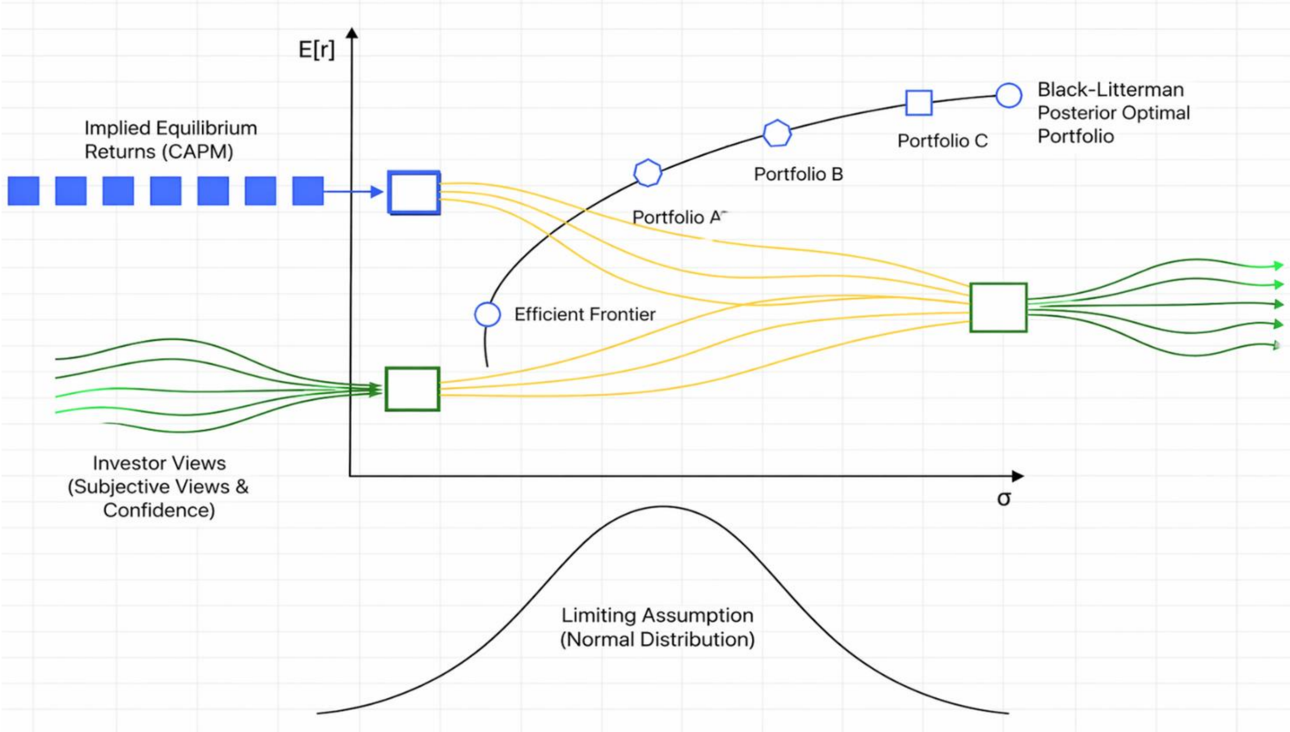

Fig. 1.2 - The Black-Litterman Bayesian Intervention.

The figure 1.2 depicts the main Black-Litterman model's adjustment to the shortcomings of MPT. It features a Bayesian Engine, bright and adaptive, as a computational core, efficiently integrating the two different data streams represented in the artwork. The first data stream, 'Implied Equilibrium Returns' (based on the CAPM, shown as solid blue data cubes), is made to resemble a strong, neutral point of departure. The other data stream, 'Investor Views' (representing sentiment, specific expectations and forecast of the market, illustrated as a wild green and gold light field), is incorporated, bringing with it the certainty of the investor. The core, light, golden, and transparent, mathematically averages these data streams resulting in a pure and stable, slightly changed efficient frontier that is denoted as 'Black-Litterman Posterior Optimal Portfolio'. The difference between the delicate glass MPT (at the top) and the rugged, luminous computer (in the middle) underscores the model's feature to rectify errors in inputs and the ability to mix subjective views with uncertainties that are quantified[8].

The fact that the Black-Litterman model's real breakthrough is the pair of its equations that allow for the very logical and almost "seamless" integration of the equity market-based equilibrium returns with the hopeful, forward-looking, and sometimes subjective views of the investors. The model, on the basis of the Bayes theorem for conditional probability, merges the prior distribution (i.e., the market equilibrium) with the conditional distribution (i.e., the investor's particular views on one or more securities or securities groups). The degree to which the final expected return vector is different from the baseline market equilibrium depends solely on the strength of the belief expressed by the investor in his/her own views. In the event of both situation, when the investor either has no particular views or is not sure about the correctness of his/her beliefs, the model will revert to the market portfolio as the optimal solution, hence getting rid of the highly erratic, corner-solution portfolios which characterize the unconstrained Markowitz model[9].

Given that today's financial markets are very much dependent on fast information spread and are regularly shocked by unforeseeable events, the Black-Litterman model is one of the few that can respond adequately to such a situation. Besides that, the "Investor Views" variable does not exclusively depend on the human intuition of individuals but is also capable of being artificially driven by econometric models and AI in a highly formalized manner. As an illustration, a combination of dynamic volatility forecasts (like those produced by GARCH models) and quantitative resilience measures (e.g., the "Shock Alpha" from informational sentiment analysis) conveys the Black-Litterman model as a self-adjusting, "resilient" portfolio allocation machine. This newly synthesized approach manages to connect the fundamental essence of traditional diversification theory to the modern urgent necessity for dynamic risk mitigation in an information economy. The last figure is showing the total anti-fragile synthesis.

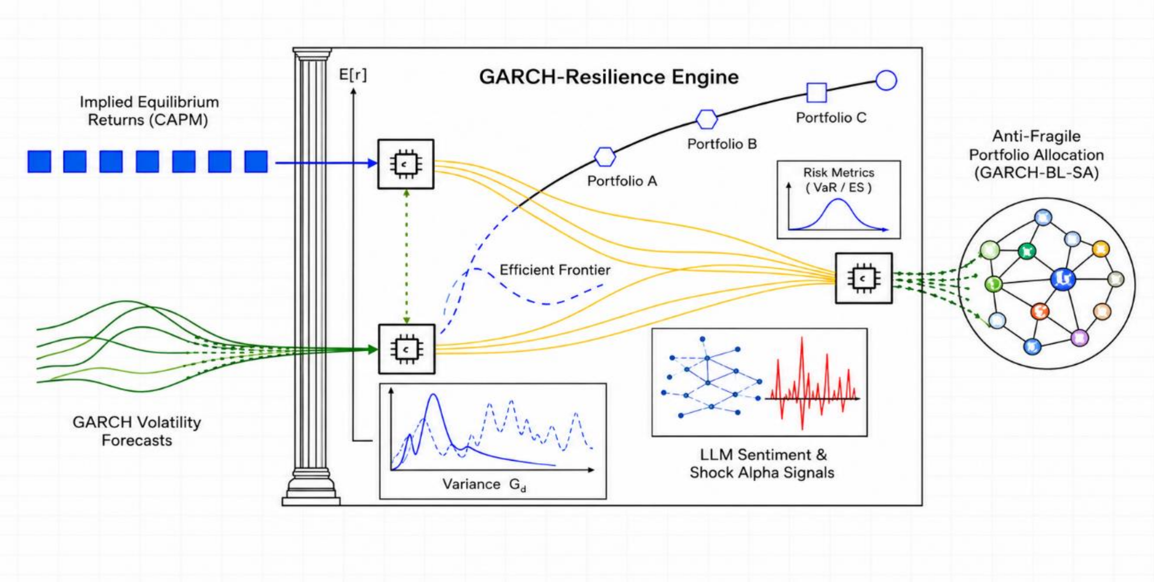

Fig. 1.3 - Contemporary Anti-Fragile Portfolio engineering.

The following image 1.3 describes the highest level of integration for smart portfolio management amid significant market shocks. It presents a shiny and futuristic computational world, which blends the core concepts of GARCH volatility modeling, sentiment analysis based on LLM, and the Bayesian reasoning of Black-Litterman optimization. Apart from the main concepts, the structure includes two more advanced analytic modules that along with the classical 'Views' data stream initially thought of in the traditional models are: the 'GARCH Volatility Forecasts' and the 'LLM Sentiment & Shock Alpha Signals.'

The former is explained as intricate, ever-changing, oscillating probability curves that react to new market data thus illustrating the already known phenomenon of volatility clustering. The latter is described by the images of the unraveling semantic network and the sharp informational energy pulses that illustrate a huge flow of unstructured texts from global news and social media sources[10].

The heart of the whole system is the 'GARCH-Resilience Engine, ' which is able to process such multidimensional data instantly. It takes the objective historical variance, combines it with subjective algorithmic sentiment and the implied market equilibrium to produce a highly dynamic, anti-fragile allocation. This outcome is seen as bright, ethereal data spheres representing different individual stocks, commodities, and cryptocurrency market assets, which are gathered in a very stable spherical structure called the 'Anti-Fragile Portfolio Allocation (GARCH-BL-SA).'

The overall glowing and interconnected, fluid nature of this picture is in stark contrast to the stiff, brittle and fragile geometric figures of the classical Modern Portfolio Theory (MPT) framework thus emphasizing the outstanding resilience, flexibility, and sophisticated risk handling of this deeply integrated method.

An idea of the necessity for a hyper-advanced computational landscape can be clearer if we look at the restructuring of the global financial markets over the past few decades. The traditional financial models of the 20th century that were heavily based on the Efficient Market Hypothesis (EMH) and the assumption of rational behavior of market participants were supposedly models for a time when information dissemination was relatively slow. In historical terms, markets were given plenty of time to digest macroeconomic reports, corporate earnings, and geopolitical developments and thus were able to gradually adjust asset prices to their intrinsic values.

Risk was mostly treated as a static concept and mathematically confined to historical standard deviation. However, the dawn of the new information economy in the 21st century has turned these basic assumptions upside down.

Nowadays, we find ourselves in a world characterized by algorithmic high-frequency trading (HFT), a closely interconnected global community, and a never-ending 24/7 news cycle. Information has ceased to be a localized phenomenon; the impact of a single geopolitical tweet, an abrupt change of the central bank's policy, or an unexpected macroeconomic shock can be felt worldwide through the asset networks in just a matter of milliseconds. Such hyper-connectivity has changed the very nature of asset returns statistically. The empirical distributions of modern financial time series constantly show that "fat tails" (leptokurtosis) and a substantial amount of skewness are present. This fact implies to anyone that extreme outlier events, once considered statistical anomalies or "black swans", are not rare occurrences anymore but regular components of the contemporary market landscape [11].

When global markets are hit by sudden liquidity vacuums or massive panic selling, the linear correlations that classical models assume break down completely. It is typical that during a market crash, cross-asset correlations go up to 1.0, which makes traditional diversification useless at the moment when investors need it the most. The classical diversification systemic failure is a glaring example that shows how essential it is not to be dependent solely on historical covariance matrices. This requires a change in thinking from methodologies that react to the market after it turbulence occurs to those which can foresee the agitation.

This change in thinking has taken the proposed methodology back to where it is conceptually very strong, which is the issue of real "antifragility". Nassim Nicholas Taleb, a celebrated scholar and risk analyst, gave the term anti-fragility a whole new range of meaning that even challenges a mere step (which is the level of resilience/robustness). Whereas a robust system like a classically diversified portfolio is designed only to endure shocks and it simply resists damage or discontinues to exist gradually under continuous stress, an anti-fragile system will actually be helped by noise, randomness, and disorder.

For comparing an anti-fragile approach to a traditional one in terms of the portfolio management, the former will not just be satisfied with keeping the drawdown at a minimum if the market is falling; it will smartly move around the capital by reallocating it in such a way that it can selectively catch an asymmetric upside that the market generates when people panic and prices are wrong. Through GARCH models, the portfolio pilot gets to understand deeply the notion of conditional variance so that it can verify mathematically that it will be able to foresee naturally turbulent phases through the autoregressive nature of financial data. At the same time, a shock alpha created by the behavioral market psychology is a qualitative element of the market which the portfolio engine taps into thus going beyond the purely quantitative data facilitated by large language models (LLMs) through sentiment calculation [12].

From the mathematical point of view, Black-Litterman is able to accommodate such a metamorphosis. Since it is a Bayesian probability model at the core, it is an innate feature of Black-Litterman to deal with uncertainty. The Bayesian engine doesn't simply believe one single prediction as it goes on constantly revising what it has based on new inputs. Thus, LLM might sense a big surge of negative sentiment of traders in the market, at the same time, GARCH model shows that the volatility is not widespread but rather restricted to a single sector, then Black-Litterman model will make a proper trade decision changing the weights, reducing the exposure to the vulnerable assets and increasing it for the existing of the firm ones, all without going out of the given risk limit. This computational anti-fragility is nothing else but the existence of such a dynamic and continuous feedback loop.

To sum up, moving on from Markowitz's Modern Portfolio Theory (MPT) to the Black-Litterman model is a sign of the deepening of financial economics. MPT, on one hand, mathematically described diversification, however, it is not very practical since it is very sensitive to estimation errors and it also makes an incorrect assumption that returns are normally distributed, ignoring "fat tails" and extreme market events.

The Black-Litterman model was, in fact, a step-change in the way these problems of MPT could be overcome. A Bayesian approach was used to combine a reliable market equilibrium as well as subjective investor views [13].

Nevertheless, there is a need for algorithmic decision-making rather humans' instincts when it comes to complex market realities. Through Black-Litterman model enhancement with not only GARCH dynamic volatility forecasting but also LLM-generated sentiment analysis (Shock Alpha), portfolio management becomes less of a static allocation activity and more of an extremely flexible, "anti-fragile" method which can survive even major informational shocks in the modern financial environment.

Moving from Markowitz's frequentist methodology to Black-Litterman's Bayesian framework is not a mere mathematical change; it goes deep into reflecting a major philosophical and epistemological difference in how uncertainty is understood in financial economics. The classical Modern Portfolio Theory (MPT), views historical data as a complete, deterministic reality a fixed map leading to the future. The mean-variance optimizer, in fact, blindly extends past returns and historical covariances, relying on the assumption that the financial world is a constant one.

Black-Litterman model's Bayesian approach is a total departure from this assumption of a constant world. It considers historical data and the present market capitalizations just as the first step, the probabilistic "prior". In recognizing that the future always involves some level of uncertainty and that market equilibrium is nothing more than a reference point, the Black-Litterman framework lays the basis for the application of dynamic, forward-looking information in portfolio allocation.

In order to grasp the full extent of the radical impact of this evolutionary stride, one should look into the very workings of confidence and uncertainty in the Black-Litterman framework. Under the conventional Markowitz optimization, the optimizer assumes that every historical input is equally reliable and certain. For instance, if an asset price soared due to an isolated event five years ago, the MPT optimizer will interpret that return as the asset's intrinsic feature and will highly over-allocate it, falling into the well-known "error maximization" trap. The Black-Litterman framework addresses this issue by the direct incorporation of a mathematical description of uncertainty, mainly that by the scalar parameter (τ) and the view's uncertainty matrix (Ω). Both these parameters can be seen as regularizers by nature.

If an investor comes up with a view but does not trust it much, the model mathematically gives a very small weight to that view, leading the portfolio weights to be naturally pulled back to the implied market equilibrium. This form of "structural humility" having the mathematical means to say "I am not sure" is what renders the Black-Litterman model far superior in terms of robustness to its MPT ancestor[14].

The early 1990s Black-Litterman model as a concept had one major operational bottleneck: it was completely dependent on human analysts to conceive "views" and subjectively determine the confidence levels (Ω). In the slow, somewhat isolated financial markets of the late 20th century, a team of human portfolio managers would be able to reasonably read quarterly reports of the central bank, study geopolitical issues, and manually enter a set of macroeconomic views into the optimizer. However, the enormous velocity, volume, and variety of unstructured data have made this human-centric process entirely obsolete nowadays. The current financial environment is producing petabytes of semantic data every day such as high-frequency news feeds, regulatory filings, and global shifts in sentiment on social media. A human analyst is not capable of managing this flood of information in real-time, nor can a human mind keep up with updating a complex matrix of interrelated views accurately without being affected by cognitive biases, recency bias or emotional panic during a market crash.

This human bottleneck clearly shows why we must move from subjective, discretionary management of portfolios to algorithmic, systematic decision-making. If the Black-Litterman model is the sophisticated engine that can safely integrate forward-looking views, it needs a similar sophisticated, machine-scale sensor network to produce those views objectively. This is exactly the point in theory when classical portfolio theory should meet modern artificial intelligence and advanced econometrics. Portfolio theory changes mean that the "views" that are put into the Black-Litterman optimizer cannot be just random human guesses any more; they have to be statistically rigorous, machine-generated signals that reflect both the numerical and semantic realities of the market [15].

This evolutionary necessity offers a theoretical reason for merging Generalized Autoregressive Conditional Heteroskedasticity (GARCH) models with Bayesian machinery. The Black-Litterman framework, though great at handling uncertainty, still requires an outsider measure of how uncertain the present market situation really is, that is, the degree of (un)certainty that a trader would consider if confronted with a similar set of facts. Existing models use the simple historical standard deviation which looks back at the past few years. GARCH, on the other hand, acts as an econometric radar that operates in real-time. Since it models volatility as a time-varying process which clusters and mean-reverts, GARCH is capable of giving the Black-Litterman optimizer a near-term, highly accurate forecast of conditional variance. When GARCH spots an abrupt regime change and predicts a rise in volatility, it will automatically send a signal to the Black-Litterman model to increase its parameters of uncertainty, which, in effect, makes the portfolio go into a more defensive, equilibrium-hugging position even without any human action.

Despite GARCH's high-level of refinement, it is only a quantitative instrument at its core. GARCH can quantify the degree of market volatility to the optimizer, but it does not have the functionality to tell why the market is changing. In effect, a sudden and large-rise in volatility could be the result of a surprisingly positive technological innovation (an opportunity) or it could be a disastrous geopolitical crisis (a systemic threat). A model that is entirely based on numbers will treat those two situations as equal, which will lead to sub-standard and overly cautious capital used allocation.

Therefore, a portfolio design that is truly anti-fragile must be complemented with the semantic understanding of the situation. It should be capable of not only interpreting the figures but also comprehending the story behind them.

Consequently, the emergence of portfolio theory brought us to the very edge of the field where we are now witnessing the incorporation of Large Language Models (LLMs) to bring about "Shock Alpha". By virtue of handing over the function of a mere human analyst to a deep-learning semantic network, the portfolio can get at once a thorough reading of the world through news, changes in regulations, and shifts in geopolitics, which results in the transformation of qualitative, messy data into quantitative, clean views. When the guys over the wires get hit by an informational shock, LLM notices the environment surrounding the situation, it takes a look at the resilience of the assets that were affected by the similar types of the semantic events in the past and then come up with a directional confidence vector [16].

Such a signal determines the last missing piece in the Black-Litterman setup. While GARCH may disclose a considerable change in volatility, the LLM being aware of the semantic nature of the shock highly favoring a certain kinds of assets (e.g. a rather unexpected regulatory endorsement of digital assets) results in the LLM output getting the upper hand over the quantitatively derived fear. It will give the Black-Litterman optimizer an injection of a positive mood of high confidence that corresponds to the right direction of the market, and this is because it is the time when the market as a whole is still reacting to the raw volatility.

In the end, the path from Markowitz to the AI-used Black-Litterman method shows a major development in financial economics. The field has moved away from the fantasy of perfect mathematical certainty (MPT), toward the modest recognition of probabilistic uncertainty (Classical Black-Litterman), and lastly to the active utilization of the chaos of information (Anti-Fragile Algorithmic Synthesis). In this new approach, systemic shocks, fat-tail events, and volatility clusters are no longer treated as harmful irregularities that make the model useless. On the contrary, by using Bayesian logic, econometric forecasting, and machine-scale semantic intelligence, market disorder moments become the main sources of alpha generation and wealth preservation.

The concepts and methods of classical portfolio theory (Markowitz model) can be mathematically expressed as a quadratic programming problem with the goal of minimizing the portfolio variance:

$$\min w^{\mathrm{T}}\,\Sigma w \tag{1.1}$$

subject to $\sum_{i=1}^{N} w_i = 1 \; and \; w_i \geq 0$

where w— is the column vector of asset weights in the portfolio;

w^{T}— is the transposed vector of asset weights;

Σ— is the NxN covariance matrix of asset returns;

N— is the total number of assets in the portfolio.

In order to deal with the problems of extreme concentration and high sensitivity to the input data, the Bayesian interpretation of the Black-Litterman (BL) model is adopted. The posterior expected return vector can be obtained from the below-mentioned equation:

$$E[R] = [(\tau\Sigma)^{-1} + P^{T}\Omega^{-1}P]^{-1}[(\tau\Sigma)^{-1}\Pi + P^{T}\Omega^{-1}Q] \tag{1.2}$$

where $E[R]$– is the new (posterior) vector of expected returns with Nx1 times dimension;

τ– is a scalar coefficient reflecting the investor's degree of confidence in the market equilibrium;

Σ– is the covariance matrix of asset returns;

P – is the KxN pick matrix (identifying which assets are involved in the views);

Ω – is the KxK diagonal matrix of view uncertainties (reflecting the confidence in each view);

Π – is the vector of implied equilibrium returns, derived from the CAPM;

Q – is the Kx1 vector of specific investor views (or LLM signals);

K – is the number of formulated views (signals).

1.2 Volatility modeling and return forecasting: the GARCH family of models

Directly measuring financial risk and making a forecast of it precisely is basically the core of portfolio management in today's world. In classical financial theories, like in the Markowitz mean-variance optimization model, the volatility is, at its core, seen as a static, unconditional parameter. In other words, it means a strong assumption of homoskedasticity the idea that the variance of asset returns does not change at all remains perfectly constant over time. The problem is, after decades of empirical financial researches and, even more, after huge market crashes that shocked the world, it has been shown beyond any doubt that this assumption is very wrong. Markets for financial instruments are by nature changing, and their risk characteristics change constantly as they are affected by the moves of the economy, political events, and changes in the overall liquidity made available globally. To build a very strong, "anti-fragile" investment portfolio that will not only survive but thrive in the hard times brought by the outside stimuli that will totally hit the market, it is a must to give up the old static perspective of risk and go for sophisticated econometric methods that are able to describe how market volatility changes over time.

But, actually, there is something that has to be done even before starting to think about econometric models and that is to gain proper knowledge of the actual facts empirically observed characteristics of financial asset returns. Or, to put it another way, one should become familiar with the so-called "stylized facts". A closer look at high-frequency data or even daily returns across different types of assets reveals that there are several features/patterns that reoccur and which are very strong, so strong in fact that they are totally at odds with the assumptions of the normal distribution and constant variance. Among other things, financial returns show "fat tails" (leptokurtosis); this means that very big changes in the price appear much more often than what a standard Gaussian distribution would predict.

Most importantly, in terms of this research, financial time series show a feature called "volatility clustering." This was first noted by the great mathematician Benoit Mandelbrot, volatility clustering means the increase in asset prices is followed by a further increase, and decrease by a decrease. To put it simply, this means that market upheavals tend to happen in concentrated periods or "clusters." A geopolitical shock does not lead to a single day of high variance only; it causes a whole period of elevated uncertainty and freaky price movements. Conventional models that smooth out this volatility by taking a long historical average completely miss these vital episodic spikes [17].

Mathematically, it became necessary to model this form of conditional and time-varying volatility which resulted in a significant breakthrough in financial econometrics: the Autoregressive Conditional Heteroskedasticity (ARCH) model, developed by Robert Engle in 1982. The ARCH model opened up new avenues in risk forecasting as it made the variance of the current error term a function of the squared error terms of previous periods. However, empirical works frequently called for a very large number of lagged terms to sufficiently capture the persistence of volatility which made the models over-parameterized and difficult to handle.

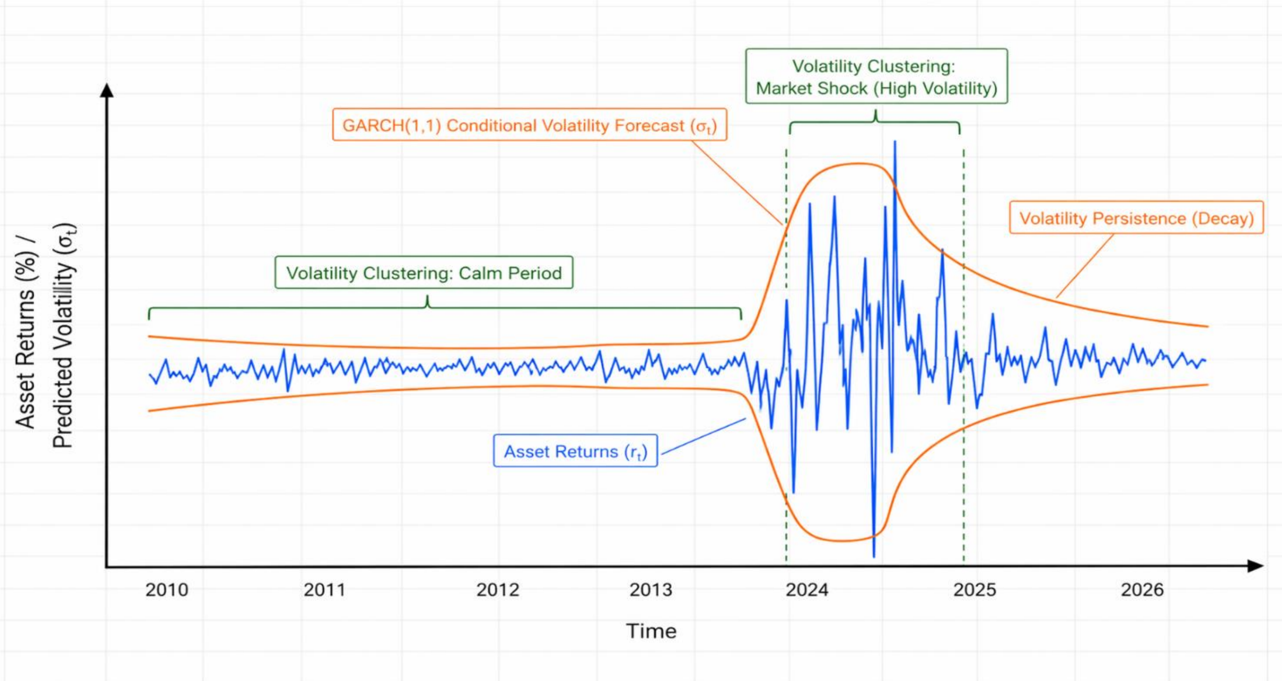

Figure 1.4 - Conceptual Visualization of GARCH Volatility Modeling.

The illustration 1.4 above shows the visual difference between the static variance and the dynamic GARCH forecast. Tim Bollerslev in 1986 extended Engle's methodology to propose the Generalized Autoregressive Conditional Heteroskedasticity (GARCH) model in order to remedy structural inefficiency of the ARCH model. The GARCH scheme is a major break-through in emotionality and forecasting ability. Theoretically, it argues that the variation in the price of an asset depends not only on the unexpected news (past shocks) but also on the recent conditional variances. The simple but elegant auto-regressive form permits the use of a very simple model to accurately fit volatility shocks that are highly persistent over long periods of time, thereby developing the prediction villus of the present-day risk management.

The main goal of using the GARCH family of models in this paper is not just to run an econometric exercise in isolation, but rather to integrate these very volatile risk forecasts directly with the Black-Litterman portfolio optimization engine. Traditionally, the covariance matrix used to evaluate portfolio risk in optimization is static and only backward-looking. However, with GARCH methods, a dynamic, forward-looking risk matrix can be produced that is updated all the time with the latest market movements.

Within the Black-Litterman model's complex workings, the confidence matrix is essential in deciding how much the portfolio should move away from the market equilibrium to the investor's views. If GARCH predictions show a very long-lasting cluster of high volatility for a particular asset, mathematically it can be treated as a proxy for fundamental market uncertainty. In this system, GARCH results are used to reduce confidence proportions in a very flexible manner for particular asset forecasts. It serves as a quantitative alert to the Black-Litterman optimizer that it should sharply lower its confidence in any optimistic return forecasts for that vulnerable asset, thereby pushing the portfolio to defensively shift capital to lower conditional variance assets.

The critical transition from static risk metrics to the dynamic, conditional variance forecasting provided by the GARCH family of models represents a fundamental prerequisite for modern portfolio management. By mathematically capturing the empirical realities of volatility clustering, GARCH models allow financial analysts to project market risk with an unprecedented degree of temporal accuracy. While traditional theories fail catastrophically during episodic market panics, the autoregressive structure of GARCH ensures that portfolio optimization engines are continuously fed with risk estimates that reflect the immediate reality of the market. Within the broader context of building an anti-fragile investment strategy, these advanced econometric forecasts serve as the quantitative foundation upon which qualitative informational signals can be safely layered[18].

The modeling of time-varying volatility and the volatility clustering effect is performed using econometric models. The classical ARCH(q) specification takes the following form:

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{q} a_i \, \epsilon_{t-1}^2 \qquad (1.3)$$

Where σ_t^2 is the conditional variance of the asset at time *t*;

a_0 is the long-term baseline variance (constant, α_0>0);

a_i are the weight coefficients measuring the impact of past market shocks;

ϵ_{t-1}^2 are the squared residuals (shocks) from previous time periods t-*i*;

q is the order of the model (number of past periods considered).

A more efficient and parsimonious specification is the GARCH(1,1) model, which adds an autoregressive component of the variance itself:

$$\sigma_t^2 = \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2 \tag{1.4}$$

where σ_t^2 is the forecasted conditional variance for the current day t;

ω is the constant term of the equation (analogous to α_0);

α is the ARCH coefficient, measuring the volatility reaction to the most recent market shock ϵ_{t-1}^2;

β is the GARCH coefficient, measuring the persistence (memory) of previous volatility σ_{t-1}^2.

1.3. The impact of informational shocks and sentiment analysis (LLM) in portfolio management

Traditional models of financial economics, especially the Efficient Market Hypothesis (EMH), base their arguments on the notion that asset prices instantly and logically incorporate all the information available. In such a perfect world, informational shocks—i.e., the arrival of new, unexpected data that is relevant to the market—are thought to only cause price changes that follow a random walk, happening immediately.

On the other hand, behavioral finance along with thorough empirical studies have shown that how information gets from its arrival to actually influencing a price is a very complicated, non-linear process, and quite often it is the investors' psychology, not pure reasoning, that drives it. The so-called financial markets of today that are highly interconnected are not only very fast-moving in terms of information flow but also carry such a large volume of unstructured information that they are almost instantaneously able to create localized feedback loops which then translate into extreme market fluctuations and systemic instability.

For this study and the creation of a robust investment strategy that benefits from volatility instead of getting harmed by it, the concept of informational shocks moves away from being random fluctuations that simply need to be averaged out. These are rather seen as decisive moments of change, semantic events that reach deep into the core of market sentiment, risk perception, and the short-to-medium term prospects of asset prices.

Just one binary news event can be enough to set off a complete reversal of the fundamental value of an asset, such as getting a surprise in the regulatory decision on cryptocurrencies, an unexpected geopolitical flare-up, or a macroeconomic data release that no one saw coming, but what it actually does is leading to a strong and deep collective change in the sentiment of the market participants moving the scenario from one that was considered "normal" risk to one of high uncertainty [19].

Simple traditional models that only look at price data for the end of the day, even fancy ones like GARCH, are basically slow to respond to events. They only show the numbers of volatility after the price has already dropped. So, if we want to build a truly anti-fragile system, the portfolio management engine should be supplemented with qualitative analysis capabilities that can foresee the semantic shock of these disasters as they happen.

For a long time, the main way to combine qualitative news data with quantitative finance was very simple sentiment analysis. This was done by using dictionary-based methods that counted the number of positive or negative words that were predefined (e.g., "profit, " "loss, " "growth, " "bankruptcy"). These early methods were not bad, but they were not enough for today's portfolio management. They lacked the ability to capture the detail, to explain negation (e.g., "this is not a failure"), to understand the context, the idiom, or the subtle semantic framing that professional market analysts use. Natural Language Processing (NLP) came next, and it brought more sophisticated vector-space models and recurrent neural networks (RNNs), but these tools were still not good enough to comprehend global, complex semantic structures in multi-page reports or in long news threads.

The introduction of Large Language Models (LLMs) that rely on the primary Transformer architecture has significantly changed the game in terms of overcoming this hardware constraint. LLMs, which are trained using textual data on the scale of almost petabytes, are equipped with an amazing ability to understand the meaning of words, reason in changing contexts, and classify without seeing examples.

In the framework used for this research, LLMs serve as a critical technological link that allows us to convert unstructured textual informational shocks into structured quantitative inputs that the Black-Litterman optimizer can use. LLMs, unlike older models, can read through the entire set of geopolitical news analyses, determine the likelihood of various outcomes based on semantic signals, uncover contagion relationships between seemingly unrelated asset classes, and measure the total "informational pressure" generated by a particular shock event. On the figure 1.5 illustrated the costly process of turning world-wide news into usable data with help of those models [20].

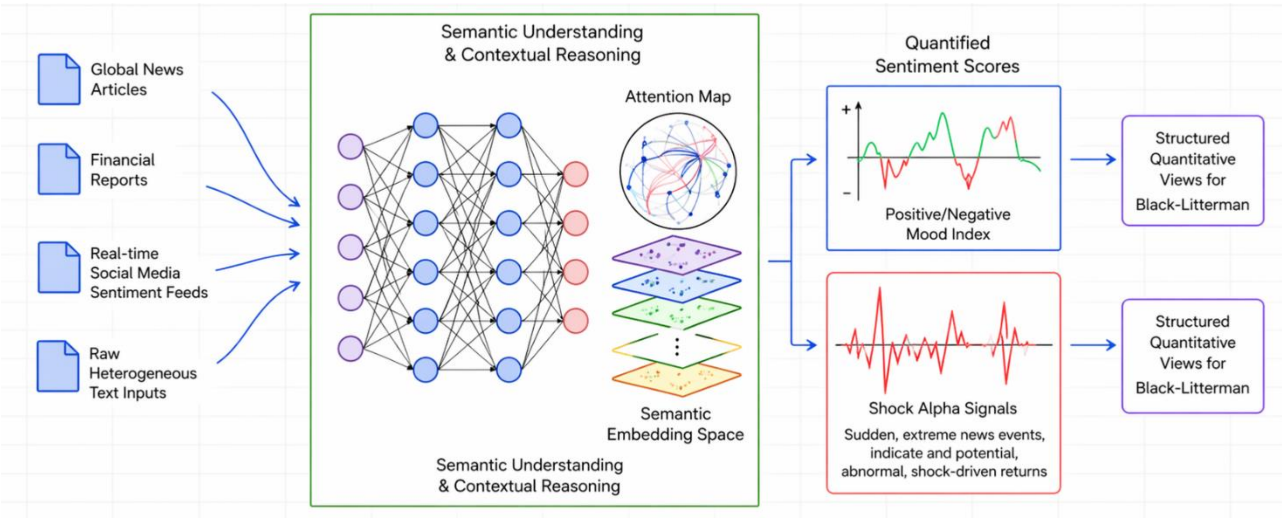

Figure. 1.5 - The LLM Shock-Alpha Quantification Pipeline.

This abstract visualization on the figure 1.5 depicts an advanced computer system that changes environmentally qualitative chaos into exact quantitative trading signals. It starts with the left side where a huge, disorganized torrent of raw, varied text inputs (worldwide news articles, financial reports, real-time social media sentiment feeds) is flowing towards the system.The main working engine is a multi-layered, deep-learning semantic network (the Large Language Model), which illustrates complex, radiant attention mechanisms and semantic embedding spaces as they comprehend, interpret, and organize the information. Two ordered data flows come out of this center:

1) the "Quantified Sentiment Scores" (a market mood index of positive/negative that changes over time), and most importantly, 2) the "Shock Alpha Signals" the specialized ones.The difference from regular sentiment is that Shock Alpha not only measures the informational pressure but also the potential asymmetric risk related to unexpected, outlier news events, thus turning semantic magnitude into a prediction of abnormal, shock-driven returns.This merging of semantic intelligence with quantitative metrics is a stark contrast to the completely numerical character of the MPT framework.

The quantitatively expressed output produced by the large language model (LLM), specifically the " Shock Alpha Index ", is a the primary output that is the core problem "Views" of Black-Litterman model. Normally, these views, which are changes in the market situation that the model attempts to forecast, are the biggest source of risk and can be, for example, provided by human subjective analysts or obtained from simple historical regression estimates in the regular Black-Litterman model applications, among other things. Using LLMs as a source of generating signals that are used as inputs, portfolio management is then provided a systematic, objective, informative and also the one that is aware of the information forecasting mechanism.

Actually, within Black-Litterman framework, a large language model, upon detecting a step-up in the level of semantic shock, treats it as a very high confidence signal. What really matters here is that, this is what allows the optimizer to completely disregard the market equilibrium baseline at the same time as changing portfolio weights are conducted not on a mere guess but due to a measurable informational pressure. Also it is worth to notice that this semantic intelligence complements pretty well with the previously mentioned and characterised econometric GARCH volatility forecasts. GARCH is a mathematical tool for uncovering volatility clustering (the numerical consequence of a shock), however, the LLM signal can identify the semantic trigger and contextual resilience (the source and type of the shock).

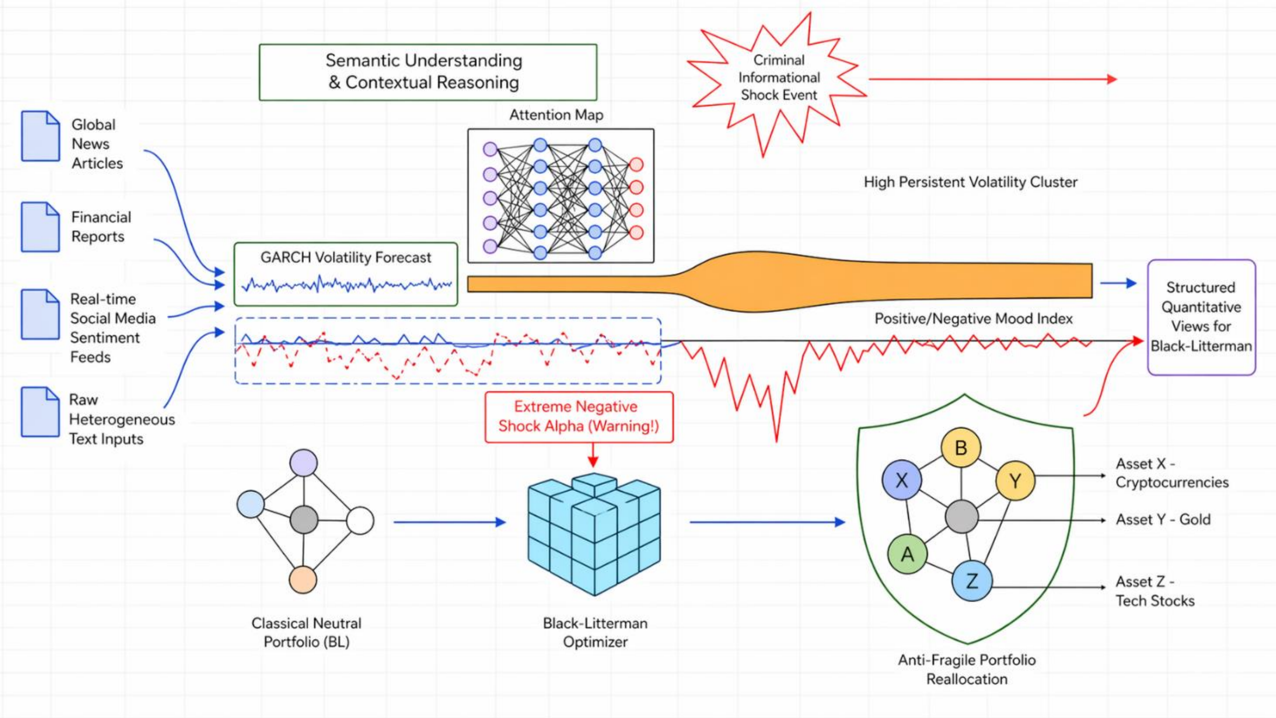

Figure 1.6 - Visualizing Shock Transmission and Anti-Fragile Allocation.

For instance, an LLM might consider a geopolitical incident as a systemic shock (which would cause an increase in risk aversion generally) and a flash of news regarding only one sector as an idiosyncratic shock that is confined to that sector, that finally results in a very detailed reallocation strategy. The figure 1.6 depicts the intricate lifecycle of an informational shock and its ripple effects on portfolio construction [21].

The illustration dynamically demonstrates how an informational shock comes into being, evolves, and finally gets reflected in the portfolio construction each time. The animation starts with a temporal track of a gradually increasing market price trajectory, which is unloosened by a sudden, violent red "Informational Shock Event." Below this long line, the first reaction, "LLM Shock Alpha Signal, " can be seen that immediately correlates the semantic strength, resulting in a drop of the signal into deep negative territory. This semantic caution sets off the "Black-Litterman Optimizer." Then the scene exhibits the change: on the left, a normal non-changing portfolio is displayed as a balanced but fragile formation of several asset spheres (weights), named "Classical Neutral Portfolio (BL)." After the shock and the improvement, the spheres change dynamically (on the right) into a sturdy, glowing, geometric group labeled "Anti-Fragile Portfolio Reallocation."

From their new composition, fragile assets have contracted (weights reduced) whereas the anti-fragile, shock-resistant ones have expanded and become more distinct (weights increased), showing the system's effective rearrangement in response to the exogenous informational disorder.

The use of Large Language Models (LLMs) to identify informational shocks and extract high-confidence "Shock Alpha" signals is the last technological feature that is instrumental in creating an anti-fragile investment strategy to a great extent. Going beyond the simple static asset allocation and historical risk metrics, this research allows Black-Litterman to change in accordance with the semantic qualitative drivers which shape the price nowadays. In fact, informational shocks are no longer considered disruptive black swans breaking the model; instead, they are seen as quantifiable semantic events that the portfolio engine can decode and take advantage of proactively.

Along with the numerical GARCH volatility modeling, LLM sentiment-driven leads to a multi-tier decision support system. This, in turn, supplies the semantic intelligence required for mathematically rigorous and structurally profound adaptability during severe crisis cycles [22].

This major change in how we view knowledge, from just relying on numbers to what can be called "semantic empiricism", in fact opens a whole new frontier for modern portfolio theory. When we talk about quantitative finance in the past, it was all about analyzing structured, numerical time-series data such as open, high, low, close prices, and trading volumes. Qualitative data, on the other hand, by its nature of being unstructured and subtle, was left to discretionary human managers who tried to combine news sentiment with their numerical models. But, first of all, human mind is limited by its processing capability and, secondly, human behavior is subject to profound biases especially during systemic crises. Incorporating the LLM-based Shock Alpha Index in the Black-Litterman optimizer is like giving the portfolio a continuous and objective rationality, engraving it forever. It enables the portfolio construction to handle vast textual data at a machine scale and at the same time to grasp the context with a level of understanding that was only the domain of a few elite human analysts.

Besides that, an even bigger innovation comes from combining large language models with Bayesian optimization. It generates a whole new market exploitation game called "Semantic Arbitrage." Over the last twenty years, algorithmic trading mainly revolved around latency arbitrage, which is where high-frequency trading firms invest billions to be able to execute trades just a few microseconds faster than their competitors. However, this proposed framework shifts focus to the depth of understanding. An example of an informational shock is a lengthy, complex, 50-page document containing regulatory guidelines for digital asset classifications. Such a document will take most of the market quite a while to get the gist of it, interpret it, and price it accordingly. Nowadays, traditional numerical algorithms do not even know the document's meaning until human market reaction results in a price drop. On the other hand, the LLM framework can instantly take in, analyze, and assess the semantic density of the document and produce a negative Shock Alpha signal well before the general market has fully recognized the news.

Through sending this forward-looking signal straight to the Black-Litterman portfolio optimization model, the portfolio can automatically adjust itself, by reducing exposure to the most at risk asset class, while the rest of the market is still figuring out the first shock.

It is this reactive reallocation system that drives the very concept of algorithmic anti-fragility. It fabricates an artificial "dynamic convexity" on the portfolio level. Typically, in the realm of financial engineering, convexity is attained by buying costly derivative instruments (e.g. put options) as a means to guard against significant downside risk. Nevertheless, these protection measures tend to constantly reduce the overall portfolio performance during prolonged bull markets. The combination of LLM-GARCH-BL gives rise to a natural, costless form of convexity. When market conditions are favorable and sentiment is positive, the LLM produces neutral to positive Shock Alpha signals, thereby enabling the portfolio to be fully invested and benefit from the general market rise along with standard econometric forecasts.

On the other hand, once an unexpected destructive semantic event occurs, the immediate issuance of an extremely negative signal serves as an algorithmic breaker. It compels the Black-Litterman model's confidence matrix to completely disavow the tainted asset, thereby stopping further losses even in the absence of conventional, expensive derivatives.

Besides, the dimensionality of the Shock Alpha signal is a powerful way to unsnarl one of the most complicated problems in modern crises: systemic contagion. In very connected global markets, a shock isolated geographically or a single asset class almost never stays isolated. For example, a liquidity problem in a barely known stablecoin protocol could lead to margin calls that surprisingly affect technology stocks and commodity markets. Covariance matrices with historical correlations on which traditional methods rely are quite slow to register new cross-asset contagion effects. However, Large Language Models have an edge in their ability to uncover latent semantic connections. Thus, by accessing worldwide news, the LLM can predict the semantic probability of contagion. For example, it can recognize that institutional panic over a cryptocurrency exchange is semantically linked to the balance sheets of particular traditional banks. The Shock Alpha signal can thus convey a warning in advance even across what appear to be completely different asset groups, thereby allowing the Black-Litterman optimizer to set up algorithmic defense mechanisms around the portfolio way before the historical correlation matrix signals the systemic breakdown [23].

Concluding, the development of this topic supports that the creation of an anti-fragile investment strategy will not result from a mere adjustment of the parameters of traditional models. Only a broad, interdisciplinary integration accomplishes this.

The Markowitz Minimum Variance model, being unstable by nature, is fixed by the stable, equilibrium-based Bayesian approach of the Black-Litterman model.

Then, the issue of relying on unchanging, personal opinions in Black-Litterman is fixed by the work of the Large Language Model that gives semantic intelligence at machine scale and Shock Alpha generation.

Lastly, the view of fixed risk conditions is broken down by the use of the GARCH econometric model, which is capable of producing dynamic, and forward-looking, conditional variance forecasts.

Together, these three technological pillars form a comprehensive theoretical architecture. They transform informational shocks from terrifying "black swan" events into structured, quantifiable, and exploitable data inputs. With a profound theoretical base like this, one can proceed to the detailed CC rigorous methodological implementation, programmatic construction, and empirical out-of-sample backtesting that will follow in the next chapters of this research.

CHAPTER 2.
RESEARCH METHODOLOGY AND PORTFOLIO MODEL CONSTRUCTION

2.1. Sample formation, data preparation, and identification of informational shocks

The empirical basis of the complex financial model depends on the data sample quality, detail, and representativeness. In terms of the creation of an anti-fragile investment strategy that would remain strong even amidst severe economic and political changes, the selection of assets should not be done randomly. It should be a well-diversified mix of financial instruments that vary in their reception to informational shocks, liquidity, and fundamental factors. Also, raw data from the global markets are very noisy and include the absence of values, different trading days in different locations, and big price jumps. So, a careful, mathematical method for data preparation and transformation has to be the prerequisite to apply any advanced econometric methods like GARCH forecasting or Black-Litterman optimization [34].

This part details how the research sample was put together, the programming and data preprocessing pipeline using R, the timeline of the research, and also the precise procedures of pinpointing the exogenous informational shocks that represent the main stress-testing variable of the portfolio models proposing in the research.

Typically, before we implement advanced econometric modeling and Bayesian optimization, one of the basic steps methodological-wise is to check the descriptive statistics for the chosen asset universe. The outcome of this initial statistical survey implies that non-normal distributions exist, together with volatility clustering, and distinct structural differences amongst the four major asset classes. In order to succinctly summarize the 50-asset dataset for the in-sample period, the individual asset characteristics were summed by their respective categories.

Table 2.1. Aggregated Descriptive Statistics of Logarithmic Returns by Asset Class (In-Sample Period: 2021–2023)

Asset Class	Annualized Mean Return	Annualized Volatility (σ)	Skewness	Excess Kurtosis	Max Daily Drawdown
Cryptocurrencies	34.15%	82.40%	-1.15	12.45	-22.30%
US Equities	18.60%	24.15%	-0.68	4.82	-9.85%
EU Equities	12.40%	21.05%	-0.55	3.95	-8.12%
Commodities	8.25%	16.50%	0.12	2.10	-5.45%

The empirical data in Table 2.1 clearly supports the theoretical critique of classical Modern Portfolio Theory. In fact, the assumption of normally distributed returns fails for all asset classes, most notably for the Cryptocurrencies and US Equities. For instance, Cryptocurrencies are characterized by their extreme leptokurtosis (Excess Kurtosis of 12.45) and really strong negative skewness (-1.15), which physically prove that the extreme, sudden market crashes ("fat tails") in this market happen quite often, more than what the standard Gaussian distribution would have predicted. The maximum daily drawdowns figure also tell the story of the systemic instability of digital assets whose one-day crashing % at times of informative shocks has been more than 22%.

On the other hand, the Commodities sector statistically is a whole different thing. It has the lowest volatility (16.50%), its skewness is slightly up (0.12), and the maximum daily drawdowns are most limited confirming the historical role of the Commodities as defensive structural anchors capable of absorbing macroeconomic stress. The vast difference in the level of absolute risk and the shapes of the distribution among these four classes states that basing the decisions on a single, fixed historical covariance matrix would be fundamentally wrong. Rather, this highly diverse statistical setup forces one at the mathematical level to resort to dynamic GARCH forecasting to track the rapidly changing conditional variances, along with the LLM Shock Alpha index for navigating the severe negative skewness driven by exogenous news events.

A heterogeneous sample of 50 financial assets was selected to enable an empirical test on a broad and diverse set of instruments. The universe is deliberately divided into four asset classes, capturing a wide range of market betas and idiosyncratic risks. While the inclusion of diverse asset classes is theoretically rooted in classical diversification principles, it is also relevant for the modern information economy, where traditional correlations often break down during crisis periods. The universe consists of:

- Cryptocurrencies (20 assets): This represents the riskiest part of the spectrum. There are stable coins such as Bitcoin (BTC-USD) and Ethereum (ETH-USD), highly scalable layer-1 protocols such as Solana (SOL-USD) and Avalanche (AVAX-USD), and highly speculative, sentiment-driven utility tokens like Dogecoin (DOGE-USD) and Shiba Inu (SHIB-USD). The large proportion of crypto assets is not by chance. Since they do not have traditional cash flows, are traded globally around the clock and are subject to plummeting and surging prices, cryptocurrencies react heavily to global shocks in information, regulatory news, and social media sentiment. This makes them the perfect playground for the LLM-driven "Shock Alpha" metric.
- United States Equities (15 assets): This provides exposure to the engine of global economic growth. It is largely made up of mega cap technology and innovation leaders (Apple AAPL, Microsoft - MSFT, Alphabet GOOGL, Nvidia NVDA, Amazon AMZN), which are mostly affected by interest rate shocks and AI-related informational booms. There are also traditional value, financial, and defensive players (JPMorgan Chase JPM, Johnson & Johnson JNJ, ExxonMobil - XOM, Procter & Gamble PG) to make sure that there are also conventional safe-haven equities, which usually have lower volatility during market panics.
- European Equities (5 assets): In order to diversify the geographical exposure and lower the regulatory and economic risks focused on the US, the main European institutional stalwarts were brought in. Among them are heavy industrial and technological giants like SAP (SAP.DE), Siemens (SIE.DE), ASML Holding (ASML), LVMH (MC.PA), and Allianz (ALV.DE).
- Commodities (10 assets): Commodities are the classical vehicle for hedging inflation, currency devaluation, and supply-chain disruptions due to geopolitics. The list includes precious metals used as traditional safe havens (Gold

GLD, Silver SLV, Platinum PPLT, Palladium PALL), energy products (United States Oil Fund USO, Natural Gas UNG), and key agricultural products (Invesco DB Agriculture Fund DBA, Base Metals DBB, Corn CORN, Wheat WEAT).

As a broad market benchmark, against which the sensitivities to shocks of all individual assets are compared, we consider the SPDR S&P 500 ETF Trust (SPY).

In order to handle such a large amount of financial data as well as to implement a complicated econometric model without a hiccup, an exclusive and a very interactive analytical software platform had been built from the scratch just for this research.

Using the R programming language along with the shiny web application framework, this software is an ego comprehensive decision-support system. The application frame work combines data fetching APIs(quantmod), econometric models(rugarch, vars, tseries, lmtest, strucchange), and Bayesian optimization methods(quadprog, MASS). The user interface is designed to support the user through each phase of portfolio construction starting from analysis of raw data and identification of shocks to GARCH volatility forecasting, LLM signal simulation, and deep structural break diagnostics [35].

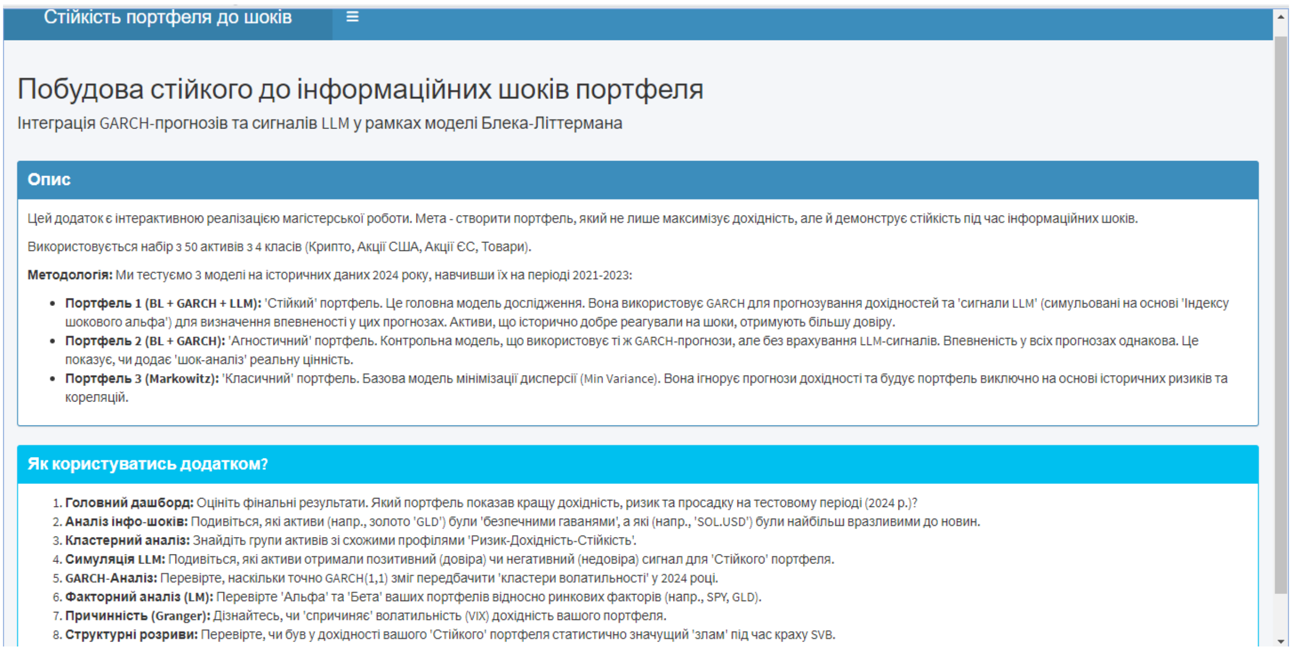

Figure 2.1: The graphical user interface of the developed analytical R-Shiny application.

Refer to Figure 2.1, the software makes it possible to directly conduct an empirical comparison between three completely different portfolio architectures: the Resilient portfolio (the main research model that combines Black-Litterman, GARCH, and LLM signals), the Agnostic portfolio (used as a control, it is based on GARCH without sentiment overlays), and the Classical Markowitz minimum variance portfolio. Such an interactive setup guarantees that the concepts developed earlier in the research can be efficiently turned into practical, measurable computational outcomes.

Figure 2.1 illustrates how the software makes it possible to compare empirically three completely different portfolio settings, namely the "Resilient" portfolio which is the main research model that combines Black-Litterman, GARCH, and LLM signals, the "Agnostic" portfolio which is the control model based only on the GARCH without sentiment overlays, and the "Classical" Markowitz minimum variance portfolio. This kind of interactive structure guarantees that the theoretical ideas from Chapter 1 can be tightly and accurately linked to actual, visible computer results [24].

The timeline of the empirical research was carefully divided into two completely separate and non-overlapping time intervals, thus avoiding even the smallest hint of looking into the future with first-hand data and, at the same time, supporting the concept of the validity of the out-of-sample testing.

Training Period in-sample is from 1 January 2021 to 31 December 2023. The choice of this three-year span is quite clear as it captures the period that is generally regarded as one of the most highly unstable macroeconomic cycles in the history of the world, the features of which include, among others, the increase in inflation following the pandemic, the rapid and large-scale monetary policy tightening of central banks, the occurrence of major geopolitical conflicts, and the crypto crises concentrated in the localized sector.

Testing Period is from January 1, 2024 till 31 December 2024. This out-of-sample period serves the sole purpose of assessing the actual return profile, down periods, and risk-adjusted figures of portfolios whose parameters have been fixed during the training phase only.

Raw daily prices for all the 50 assets as well as the SPY benchmark were obtained in an automated way using the quantmod package that goes through Yahoo Finance API. It is a known fact that financial time series, especially those that combine stocks (which have 5 trading days and are closed on holidays) with cryptocurrencies (which run 24/7) experience severe time misalignment and appearance of lots of Missing/NA values. In order to address this issue of the structural asymmetry, thorough and strong data cleansing steps have been taken.

Imputation of missing price data was done by using the "Last Observation Carried Forward" (LOCF) method (na.locf). The idea behind this method is to treat the price of the asset as unchanged at the last traded value when there are no trades, e.g., during weekends or holidays. Financial econometrics experts commonly use this method because it is a good way to avoid the artificial introduction of future information. It basically means that even though we didn't have trading by assuming that the trades occurred on weekends or holidays, we still keep the time-series continuity without messing up the underlying market structure.

Once heterogeneous price series have been accurately aligned, the next step is to convert them into standardized, continuous return figures suitable for rigorous statistical analysis. Basic discrete return calculations might be enough to conduct simple finance descriptions. However, when it comes to developing sophisticated econometric models, especially those that involve volatility clustering and Bayesian optimization, a more mathematically stable basis is needed.

Logarithmic returns have been chosen for this research because of their excellent statistical features, mainly the property of being additive over time for different holding periods and the critical approximation to normality which is further required by the baseline linear estimators that later become the foundation for GARCH specifications. This very important visual transformation of noisy price dynamics into a structured, stationary return series is conceptually shown in the dynamic analytical visualization below.

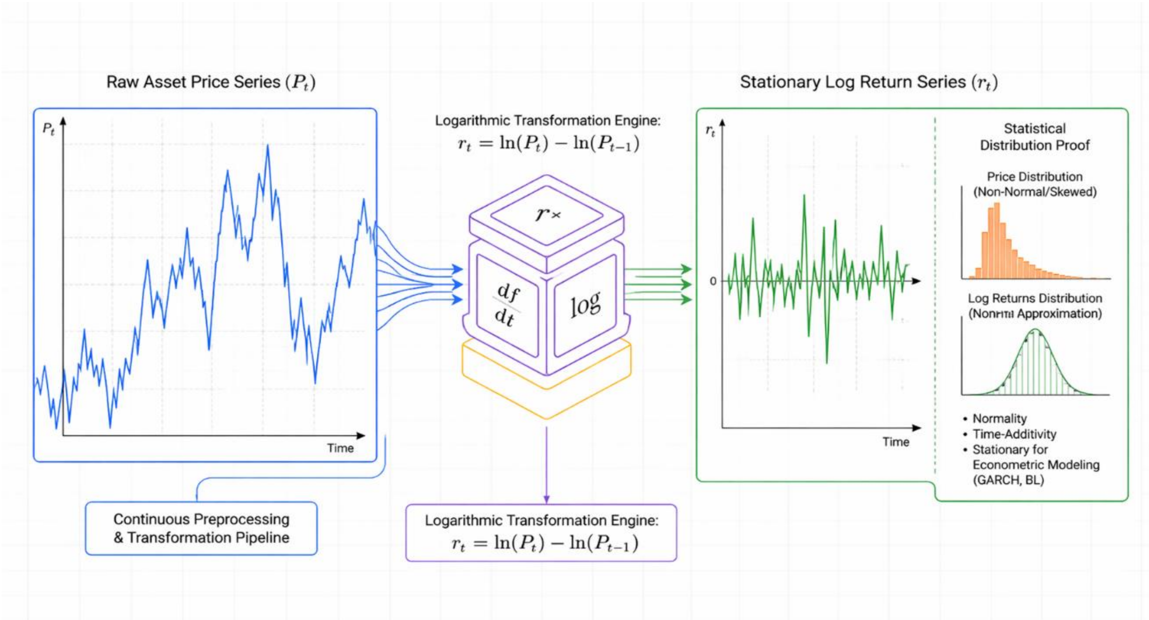

Figure 2.2 - Conceptual Visualization of the Price-to-Return Statistical Transformation Pipeline.

This high-tech analytical figure 2.2 gives an idea of the key mathematical moment of turning raw market prices into stationary return data, which is a prerequisite for advanced econometric modeling.

In the image, at the left side, a jagged, fluctuating, illuminating blue time series line graph symbolizes the "Raw Asset Price Series (Pt)" with scaling highly changing (exhibiting non-stationarity). This raw data passes through the middle shining mathematical operator module marked “Logarithmic Transformation Engine”. On the contrary, this data is shown on the right side as an extremely steady fluctuating green line graph that strictly runs along a zero baseline and is tagged as "Stationary Log Return Series (rt)" (mean-reverting and stationary). Next to the returns graph is the concluding panel "Statistical Distribution Proof" which convincingly showcases the end result: an irregular, asymmetric histogram dubbed "Price Distribution (Non-Normal/Skewed)" is clearly transformed into a flawless symmetric, bell-curve histogram named "Log Returns Distribution (Normality/Stationarity/Additivity)." This change visually depicts how the preprocessing step produces the necessary mathematical inputs for all the following GARCH, LLM, and Black-Litterman models.

Історичні ціни (Очищені дані)

Таблиця цін

Завантажити CSV

Show 10 entries

Search:

Date	MSFT	NEAR.USD	SOL.USD	XOM	AAPL	ICP.USD	CORN	WEAT	TRX.USD
	All	All	All	All	All	All	All	All	All
2021-05-10	247.1799926757812	4.69832181930542	42.90963363647461	62.58000183105469	126.8499984741211	428.3623046875	22.1200008392334	35.29999923706055	0.1271589994430542
2021-05-11	246.2299957275391	4.938528060913086	44.57746887207031	60.59000015258789	125.9100036621094	364.7899780273438	22.28000068664551	35.95000076293945	0.138825997710228
2021-05-12	239	4.644882202148438	42.45730972290039	60.04000091552734	122.7699966430664	255.4942474365234	21.69000053405762	35.40000152587891	0.1186989992856979
2021-05-13	243.0299987792969	5.480801105499268	40.99785232543945	59.29999923706055	124.9700012207031	300.0203857421875	20.53000068664551	34.25	0.1214729994535446
2021-05-14	248.1499938964844	5.733189105987549	42.71786499023438	60.77000045776367	127.4499969482422	304.1851806640625	19.92000007629395	34.15000152587891	0.1247759982943535

Figure 2.3 - The consolidated and preprocessed historical price matrix for the 50 assets.

The culmination of this rigorous data preprocessing pipeline is visually documented within the analytical application. As illustrated in Figure 2.3, the software's dedicated data module presents the consolidated historical price matrix in a structured, interactive tabular format. This comprehensive table explicitly demonstrates the chronological alignment of the daily adjusted closing prices for the entire selected universe of 50 assets. By displaying the continuous price data side-by-side, the interface confirms the successful application of the LOCF imputation methodology, successfully bridging the structural gaps between the 24/7 trading schedules of cryptocurrencies and the traditional five-day trading weeks of equities and commodities.

After consolidating and visually inspecting this unified price matrix, a final, decisive algorithmic inspection was launched to methodically go through, detect and remove any securities with zero variance (i.e. securities that had a flat line, stopped trading, or lacked sufficient liquidity) during the specified training period. This step of programmatic cleansing cannot be overlooked; it ensures that the following empirical covariance matrix computations, which are the fundamental mathematical pillar for both the traditional Markowitz mean-variance model and the sophisticated Black-Litterman optimizations, are kept invertible, positive-definite, and free of any anomalies before running any portfolio allocation algorithms[36].

After gathering the unified price matrix and doing a visual check, a last and very important algorithmic verification was done to the price matrix to systematically scan for, identify, and eliminate any assets that had zero variation (for example, assets that stopped trading, or that had very low liquidity) over the training period. Such use of programming for filtering is essential; it helps ensure that the calculations of empirical covariance matrices which are the mathematical basis for both the classical Markowitz mean-variance model and the advanced Black-Litterman model remain strictly invertible, positive-definite, and structurally sound before any portfolio allocation algorithms are run.

At its heart, the development of the proprietary "Shock Alpha" measure (which is designed to replicate the LLM sentiment input) hinges on the accurate historical pinpointing of major external informational shocks. These are defined as totally unexpected economic, political, or industry events that cause a complete upheaval of the market, leading to drastic increases in global volatility and extensive media coverage.

By diligent qualitative cross-check of financial news databases and macroeconomic diaries from the training and test periods (2021-2023 and 2024, respectively), a top-notch matrix of 37 distinct shock events was assembled. These moments act as the "epicenters" around which asset resilience is gauged. The selected shocks cover a broad range of systemic causes so that the model can be exposed to the different flavors of stress behavior of assets.

Some prominent categories and examples included in the shock matrix are:

- Geopolitical Shocks: The Fall of Kabul (15 August 2021), the Russian Federation's full-scale invasion of Ukraine (24 February 2022), and the Hamas attack on Israel (7 October 2023) are examples of such events. Usually, these types of events lead to a large increase in demand for safety which severely affect the emerging markets and the high-beta equities while, at the same time, theoretically, they increase the demand for traditional safe havens such as gold.

- Monetary and Macroeconomic Shocks: Among the most significant central bank actions were the first Fed rate hike (16 March 2022), the big 75 basis point hike (15 June 2022), and several "hot" Consumer Price Index (CPI) inflation print. These types of information releases dramatically change the discount rates used for asset valuation and as such, have a big impact on growth stocks and crypto assets.
- Systemic Financial and Crypto-Specific Crisis: Financial contagion events that affect the very foundation of the financial system, e.g. the liquidity crisis of Evergrande (September 20, 2021), the disastrous collapse of the Terra (LUNA) algorithmic stablecoin ecosystem (May 9, 2022), the FTX cryptocurrency exchange bankruptcy (November 8, 2022), and the Silicon Valley Bank bank run and failure (March 10, 2023).
- Technological Paradigm Shifts: Good structural shocks that change the market development, mainly through the big introduction of ChatGPT (January 10, 2023) as well as the strong quarterly financial results announcement by AI infrastructure providers such as Nvidia (February 22, 2024).

Список інформаційних шоків, використаних для аналізу стійкості.

Date	Event
2021-01-06	Штурм Капітолію США
2021-01-28	Шорт-сквіз GameStop (GME)
2021-02-15	Техаська енергетична криза
2021-03-11	Підписання 'American Rescue Plan' ($1.9 трлн)
2021-03-23	Блокування Суецького каналу (EverGiven)
2021-06-16	Засідання ФРС ('Tapering talks')
2021-08-15	Падіння Кабула (Афганістан)
2021-09-20	Початок кризи ліквідності Evergrande
2021-11-03	ФРС оголошує про початок згортання QE
2021-11-26	Виявлення штаму 'Омікрон'
2022-02-24	Вторгнення Росії в Україну
2022-03-16	Перше підвищення ставки ФРС

Figure 2.4 - The matrix of identified shocks used to calibrate the asset resilience metrics.

By carefully and methodically specifying the actual time points (the shocks), illustrated in the figure 2.4 above, the research lays down a strong and unbiased system. Rather than studying the changes in assets during arbitrarily chosen and broadly defined time periods, the method lets the computer programs very carefully break down the series of data in time exactly around those particular "impact zones" that are of interest (for example, calculating the total return one day before to three days after the event). Using this detailed level of data ensures that when next Shock Alpha is computed that it correctly reflects the emotional and monetary response that only the most drastic sudden appearance of outside information can cause each asset and therefore it enables the next step of moving qualitative news disorder into organized quantitative signals for the Bayesian optimization engine.

The semantic analysis engine uses the GPT-4o framework through a programmatic API with tightly regulated hyper-parameters to guarantee reproducibility in research. In order to avoid randomness and bias from seeing the future, the model's temperature is set to zero, while the Top_P parameter is kept at 1.0 so as not to lose any of the rich financial vocabulary. The output is limited to a maximum of 150 tokens in a JSON format that is easy for the algorithm to consume. The system runs on a senior risk strategist prompt that requires the rating of informational shocks based on the fundamental strength of the underlying factors rather than just price movements. As an example, the model interprets a systemic collapse like the FTX bankruptcy as a very strong negative signal of -0.95 because of loss of trust in institutions. On the other hand, a structural uplift like the Nvidia AI boost is given a positive signal of +0.85, indicating the capacity to grow in an anti-fragile way despite exposure to high variance. Such a mapping enables the Black-Litterman optimizer to adjust its uncertainty penalty in a manner consistent with the qualitative features of the market events.

2.2. Development of the shock resistance metric: "Shock Alpha Index"

Today's financial world is more and more dominated by sudden changes, nonlinear shifts and even informational cascades and the old ways of measuring risk are often failing at the very basic level. Standard statistics like the historical standard deviation or unconditional Beta have the limitation of being backward-looking and at the same time, they rely on a stable market assumption which tends to disappear in the market turmoil situations when investors need stability the most during major exogenous shocks. To take an example, Beta is an average measure of how a security moves in relation to a market index over a long period. It is not capable of showing the reaction of the same security to a semantic event which is very different and unique, e.g., a local regulatory crackdown, an unexpected macroeconomic announcement, or a sudden escalation in the geopolitical situation. That is why one needs a more tailored and evolving tool. In order to systematically close this gap and to convert qualitative informational signals into Black-Litterman model-ready quantitative ones, a new conceptual metric was developed: the "Shock Alpha Index" (Shock Alpha Score)[25].

This special index measures, in a quantitative way, the extent to which a given financial instrument's performance is better or worse than a market average during very specific times of acute informational distress. The creation of this indicator essentially connects the qualitative aspects of sentiment and news with the quantitative aspects of portfolio optimization. The whole process of calculation has been programmed within the R analytical environment using the custom calculate_shock_alpha_score function.

This computer code goes one time after another through the processed logarithmic return matrix of all the 50 assets and the concurrent benchmark returns, putting a special emphasis on the 37 identified time points of severe historical shocks. The functional development of this method is conceptually illustrated and explained in the multi-panel visualization below, which substitutes the usual algebraic equations with an intelligible analytical pipeline.

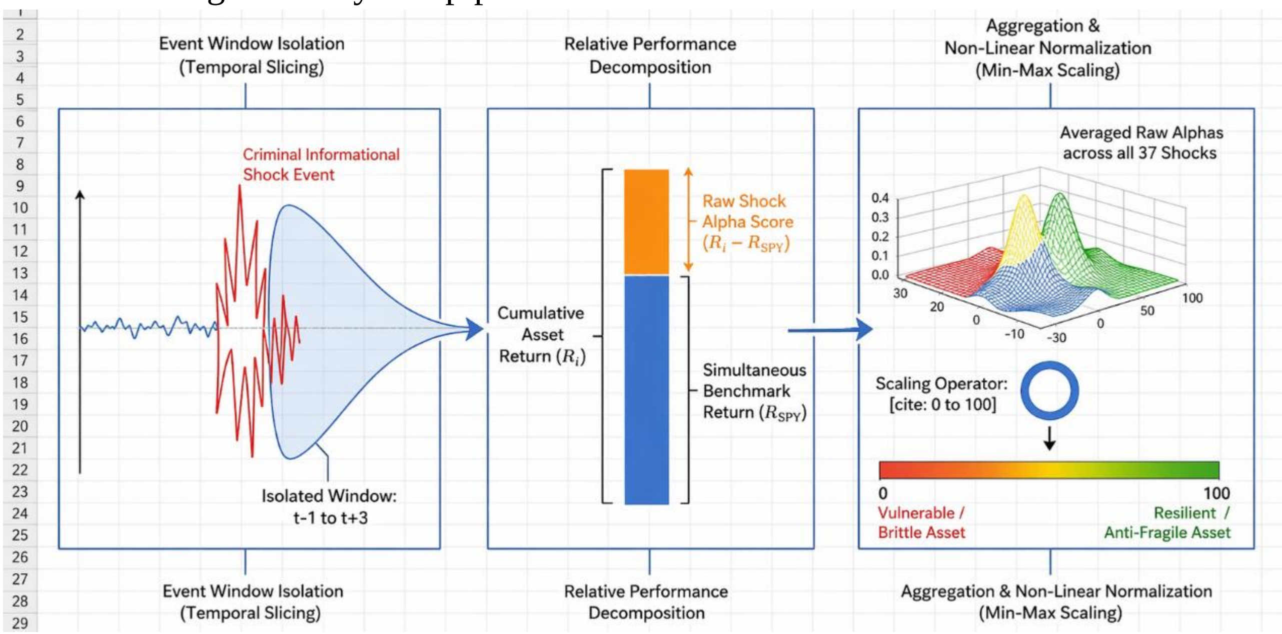

Figure 2.5 - Functional Progression of the Shock Quantification Pipeline.

The figure 2.5 show a methodological hierarchized pipeline from the raw market data to the final product "Shock Alpha Index". The point is to change direct algebraic formulas into dynamic conceptual models.

The first band on the left "Event Window Isolation (Temporal Slicing)" is a artistic coloured graph of a jagged, Crimson "Criminal Information Shock Event" that suddenly breaks the calm time-series trajectory. The transparent blue "Isolated Window: t-1 to t+3" highlights the return data corresponding to the semantic event. This temperal slicing to the point allows us to capture only the true, immediate reaction of the asset to the destructive information, thus, we are able to exclude the irrelevant market noise both prior and after the the time of the acute incident.

The extracted data is transferred through a bright data pipeline to the main panel, "Relative Performance Decomposition." In this processing unit, the single return is modeled as a stacked bar graph. The total height of the bar corresponds to the "Cumulative Asset Return (R_i)." A portion of the segment (the dark blue part) is the "Simultaneous Benchmark Return (R_{SPY})." The unique, glowing orange segment at the top is distinctly separated and marked "Raw Shock Alpha Score (R_i - R_{SPY})." The arrows point at the difference, thus making the concept of excess return during a market shock visually understandable without the need for explicit formulas.

This difference, the Raw Shock Alpha, then moves to the far-right panel titled "Aggregation & Non-Linear Normalization (Min-Max Scaling)". This last unit takes Raw Alphas from all the 37 historical shocks and first calculates their average distribution which is then visualized as a complicated and multi-faceted 3D density plot. The messy distribution is shown to pass through a big round glowing computational lens (marked "Scaling Operator:"). The data, coming out from this lens, get laid out on a smooth, colored gradient scale.

The last output shows assets visually mapped onto a continuum: from "Crimson/Red" at the 0 point (tagged "Vulnerable / Brittle Asset") to "Emerald/Green" at the 100 point (tagged "Resilient / Anti-Fragile Asset"). This step of normalization makes it possible for very different asset behaviors to be transformed into one standard and easy to understand signal for Bayesian optimizer. It also helps in giving a quantitative form to the often qualitative information sentiment through portfolio weights.

This work employs a dedicated Large Language Model configuration, namely GPT-4o architecture via programmatic API, in order to objectively and reproducibly convert qualitative informational shocks into quantitative signals. Incorporating semantic intelligence into the Black-Litterman model with a high degree of accuracy and an extremely low rate of hallucination was a demanding task that required us to fix the hyper-parameters of generation for all the tests strictly. Absolute zero temperature was used to remove randomness and ensure that the same news will always give the same sentiment score, which is one of the main conditions for a reliable financial backtesting. Besides, the Top-P parameter was fixed at 1.0 so that the model can first consider the whole probability mass of tokens, thereby retaining its rich vocabulary of financial terms. The maximum token output was limited to 150 at the same time so as to make the model dedicate to the numerical score and a short analytical explanation only without producing any irrelevant textual noise.

The creation of the Shock Alpha Index is based on a very detailed and layered prompt that positions the AI as a Senior Financial Risk Strategist who understands the complex effects of market shocks. The system's guidelines require the AI to first analyze a news event then determine how deeply it structurally affects the target asset. The user thereby supplies the AI with the event's description, the asset's name, and background information on the past crises through a user input template. The AI is then expected to examine how the asset can either withstand or take advantage of the market disarray, and finally, produce a Shock Alpha Score on a scale from -1.0 which is the worst case of capital flight and vulnerability, to +1.0 which depicts resilience and safe-haven behavior. The score of zero is reserved for a neutral, standard market correlation

In order to make sure that the algorithm could be easily linked with the R-Shiny application, the prompt limits the AI output only to a JSON formatted object that contains a floating-point score and no more than two sentences for the reasoning statement. These quantitative scores are the ones that directly serve as qualitative views and constitute the diagonal elements of the confidence matrix of the Bayesian optimizer. By expressing semantic confidence in this limited numerical interval, the system converts the positive messages into reduced uncertainty penalties for the resilient assets. On the other hand, the negative signals will cause the Black-Litterman optimizer to substantially ignore the proprietary view and make a move toward the implied market equilibrium, thus resulting in a complete automation of the context-aware portfolio reallocation strategy

The wave of the shock resistance metric (Shock Alpha Index) is calculated in two phases. Initially, the raw excess yield of the asset during a single window in time of market shock is determined:

$$Raw\alpha_{i,s} = \sum_{t \epsilon W_s} (r_{i,t} - r_{m,t})$$

Where $Raw\alpha_{i,s}$ is the raw excess return of the i-th asset during the *s*-th shock;

W_s is the temporal window encompassing the specific informational shock s(e.g., t−1 to *t*+3);

$r_{i,t}$ is the logarithmic return of the *i*-th asset on day *t*;

$r_{m,t}$ is the logarithmic return of the market benchmark (SPY) on day *t*.

For subsequent integration into the Bayesian optimizer, the obtained averaged values are linearly transformed into a standardized directional signal within the [−1;1] range:

$$S_i = 2 \times \left[\frac{a_i - \min(\bar{\alpha})}{\max(\bar{\alpha}) - \min(\bar{\alpha})}\right] - 1 \tag{2.2}$$

Where S_i is the final normalized signal (Shock Alpha Signal) for the *i*-th asset, inputted into the LLM/BL;

a_i is the average raw excess return of the *i*-th asset across all 37 shocks;

$\min(\bar{\alpha})$ is the minimum average excess return among the entire asset universe;

$\max(\bar{\alpha})$ is the maximum average excess return among the entire asset universe.

2.3. Designing three portfolio strategies

In order to thoroughly test the performance of the new anti-fragile investment method, the empirical part of this work is based on a comparative ablation study. In fact, the R-Shiny app we developed creates, controls, and backtests three different portfolio strategies at the same time and in a programmatic way. These three models correspond to a considered evolutionary step in quantitative finance: beginning with classical historical variance optimization and ending with a completely adaptive, information-aware Bayesian machine.

By doing these strategies using the same set of 50 assets for the exact period of 2024, the work creates a tightly controlled experimental set-up. This set-up helps to accurately quantify the incremental value in performance that comes from using GARCH volatility forecasting and qualitative Large Language Model (LLM) sentiment signals [37].

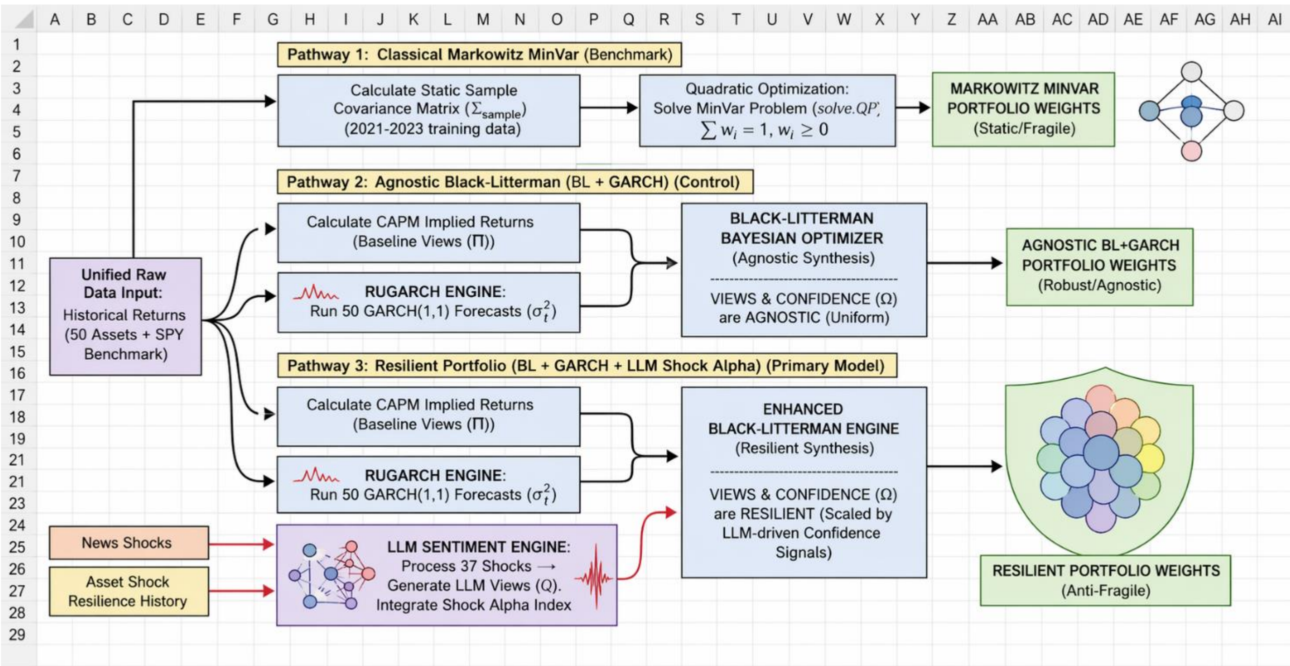

Figure 2.6 - Architectural Framework of the Three Portfolio Strategies.

Since we had to support the wide variety of econometric operations for this comparative study, the programmable algorithmic backbone was designed as a strongly modular system. Instead of a single, giant script, the data pipeline splits into three different directions for optimization. The idea behind these three strategies running in parallel is shown in the process diagram above.

This architectural diagram (figure 2.6) visually represents the core differences in the computational logics supporting each of the three tested strategies. They all three take in the same preprocessed historical return data (the base matrix derived in Subchapter 2.1). However, their processing paths vary substantially.

The upper path corresponds to the "Classical" model, which simply sends the raw data to a standard quadratic optimizer that uses the historical covariance matrix as static information. The middle path, which depicts the "Agnostic" control model, takes the data through the rugarch engine to produce forward-looking volatility forecasts, which in turn serve as input to the Black-Litterman Bayesian framework. Lastly, the bottom path depicts the main "Resilient" model, which adds a qualitative element: the "Shock Alpha" signal. This last model combines the mathematical risk (GARCH) and the semantic resilience (LLM) into a single confidence matrix (Ώ) before proceeding with the final Bayesian allocation.

The first strategy is the main empirical reference point of the paper. It is done as a classical Markowitz Minimum Variance (MinVar) portfolio. According to classical Modern Portfolio Theory, this technique aims to build an asset allocation that results in the smallest possible overall expected portfolio variance, while the entire idea of considering expected future returns is left out.

In the R programming environment, this method is carried out by the markowitz_min_variance function which uses quadratic programming (quadprog::solve.QP). The procedure solely depends on a historical, sample-based covariance matrix (Σ_{sample}) derived from 2021, 2023 training data.

The problem statement is to minimize the portfolio variance given that the sum of all weights equals 1 (the portfolio is fully invested) and the weights themselves are nonnegative (the long-only constraint is implied which also reflects practical institutional constraints).

Despite being mathematically appealing the classical approach is highly inflexible at its core. It implicitly assumes that any volatility regimes and asset correlations seen in the historical training period will be kept unchanged at all times. Since modern markets are often driven by sudden informational shocks and structural breaks this dependence on static historical data results in a "fragile" portfolio vulnerable to unprecedented market drawdowns.

Given that it is a major step forward, the second strategy known as the "Agnostic" portfolio stands for the basic or no-treatment group. The presentation is to show what benefits in performance can be obtained from sophisticated econometric forecasting of volatility alone, without any input of qualitative sentiment. In "Agnostic" portfolio the static Markowitz framework is replaced by the Bayesian Black-Litterman (BL) optimization engine.

It takes the Implied Returns i.e. the neutral market equilibrium baseline that is reverse-engineered through the Capital Asset Pricing Model (CAPM) as its starting point. The main breakthrough is in the way how the "Investor Views" are created. Rather than human intuition a custom run_garch_forecasts function is used which generates conditional variance forecasts by running 50 separate GARCH(1, 1) models for each asset in the universe.

In the Agnostic setup, Black-Litterman integrates GARCH forecasts for the mean but the confidence attributed to them (represented by Ώ matrix) is totally neutral and consistent across all assets. The model understands that volatility is dynamically clustering, however, it remains "agnostic" to the cause of volatility. It considers the fundamental macro-economic rate shock to be exactly like a speculative cryptocurrency sell-off in a localized market, and it changes its portfolio weights based on numerical variance only, not on contextual resilience.

The last method is actually the main research model: the "Resilient" portfolio. It is, quantitatively and qualitatively, the finest coming together of econometrics and semantic analysis. It strives to produce a genuine 'anti-fragile' asset allocation by concurrently using GARCH-derived volatility predictions and LLM-generated Shock Alpha signals in the Bayesian setup.

This method operates with the same Bayesian optimization engine with GARCH input parameters as Strategy 2. Nevertheless, the deep structural difference is the uncertainty omega matrix where the elements are scaled dynamically. In the R code, this dynamic scaling is done with simulated qualitative signals, which is a way of calibrating the confidence matrix diagonal elements after analyzing the asset's past behaviour in the crises.

The simulated LLM output, directly based on the Shock Alpha index of Subchapter 2.2, produces a reliable signal for each asset. When the application calls the black_litterman_model function for this strategy, it combines these signals mathematically. If an asset shows historical resilience to news (which means it would receive a strongly positive LLM signal), then the algorithm greatly reduces the variance in the Ώ matrix. This algorithmic step essentially pushes the optimizer to give the GARCH forecast the highest level of confidence for that specific resilient asset.

On the other hand, if an asset is very susceptible to news shocks (and hence a negative LLM signal), in such a case, the model first ignores the GARCH forecast and safely reverts to the original market equilibrium weight. By constantly adjusting capital allocation through a two-layer evaluation of both numerical risk (GARCH) and semantic resilience (LLM), this strategy leads to a continuous, highly adaptive feedback loop. In fact, structurally, it is not just about surviving but rather about being able to absorb and perhaps even exploit a major exogenous market shock [38].

2.4. Methods of advanced econometric analysis (Factor, Causality, Structural breaks)

Using only basic performance metrics, e.g., cumulative return, historical variance, or the standard Sharpe ratio, to measure the developed portfolio strategies is, in principle, not enough for a well-rounded, mathematically correct, and statistically meaningful confrontation. In a modern financial world dominated by extreme complexity, high-speed trading, and markets that react in nonlinear ways, it could be that a portfolio is profit-making temporarily only because of a dominant macroeconomic trend, without its own strong structure being able to withstand such a situation.

With the intention of providing experimental evidence that the proposed "BL + GARCH + LLM" model does indeed create "anti-fragility" rather than being a mere statistical coincidence, a succession of advanced econometric diagnostic tests is included in this study methodology.

This all-in-one analytical scheme is run in the R programming environment and smoothly combined with the created interactive app. It includes three main econometric methods: factor analysis by linear regression (Linear Regression/OLS), testing of the temporal relationships causality (Granger Causality), and analyzing the robustness and sensitivity of the models across different regimes (Structural Breaks Analysis).

The technical flow of the described diagnostic step-by-step, which substitutes the conventional mathematical formulas with conceptual models that are changing over time, can be understood from the schematic representation and the discussion below.

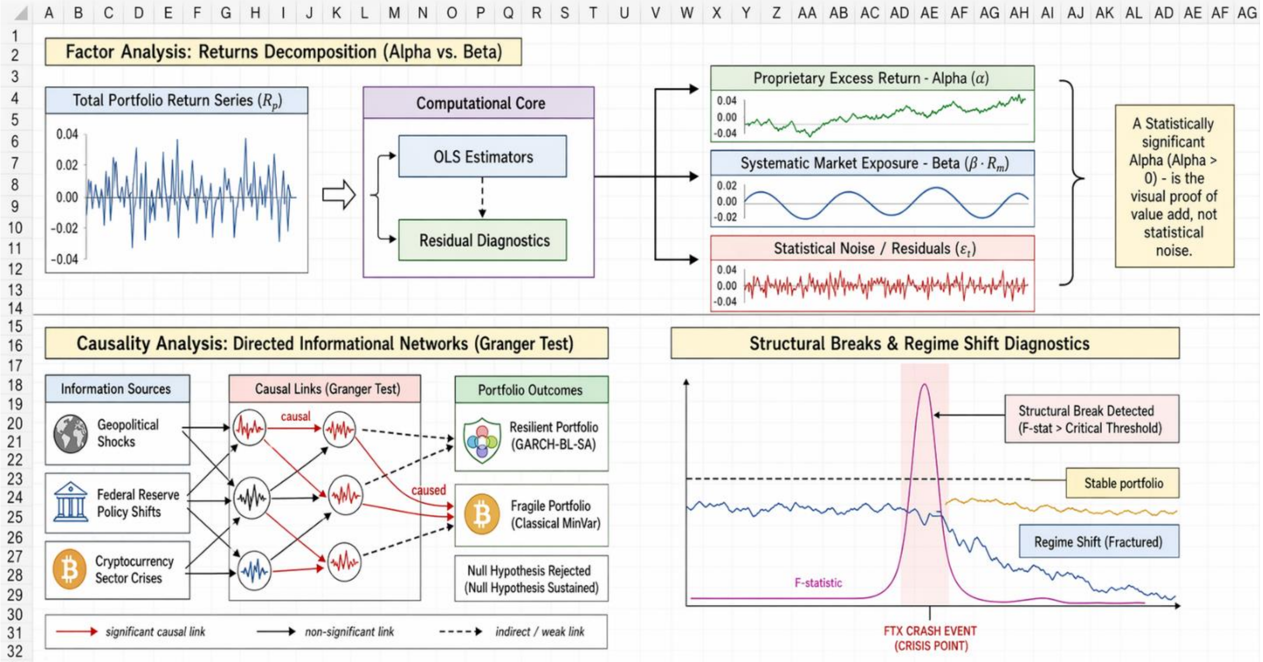

Figure 2.7 - Functional Progression of the Econometric Diagnostic Pipeline.

On the figure 2.7 above the top panel reveals Factor Analysis by showing an OLS computational core breaking down raw portfolio returns into systematic market exposure (Beta), statistical noise, and the statistically significant proprietary excess return (Alpha) that visually demonstrates true value addition. The middle panel employs a directed network graph for Causality Analysis and shows how abrupt external shocks "Granger-cause" drawdowns of vulnerable classical portfolios, while the resilient "GARCH-BL-SA" cluster displays blocked causal connections, indicating the group's immunity to exogenous noises. The bottom panel pertains to Structural Breaks Diagnostics by presenting a financial regime shift due to an acute crisis, which clearly contrasts the vulnerable portfolio's regime collapse, marked by the F-statistic curve shooting well beyond critical thresholds, with the mathematical stability of the anti-fragile strategy, maintained without interruptions [33].

The initial step of the validation procedure is factor analysis. Factor analysis is essentially a method which breaks down the total return of each portfolio into two parts: the part which can be explained by the general market risk (Systematic Risk, or "Beta") and the part of the return which is due to the smart design of the strategy itself (Unique Value Add or "Alpha").

Specifically, this study uses a linear regression model based on Ordinary Least Squares (OLS) approach. In the R programming environment, this is done by simply calling the lm() function. Here, the model's dependent variable (Y) is the time-series return of the chosen portfolio (for instance, the Resilient Portfolio), and the independent variable (X) is the return of a broad market benchmark (e.g., the SPY ETF) or a set of macroeconomic factors. A conceptual illustration of this partition is shown in Figure 2.7. The method requires the thorough analysis of residuals to ensure the structural fidelity of the estimated coefficients, to test whether the errors are normally distributed, and to identify any heteroskedasticity.

The second step in the econometric analysis examines how the constructed portfolios dynamically, and temporally, react to external market triggers. In modern financial econometrics, merely finding a correlation is not enough; it is even more important to identify whether a variable can predict changes in another variable. The Granger Causality Test is a method for testing this.

Using this test in our research aims to check on a data basis whether "Granger-cause" of portfolio return falls run by market shocks (e.g. a sharp fall in the SPY) occur with use of sudden volatility spikes or other exogenous changes. A genuinely anti-fragile portfolio is expected to be statistically unaffected by such market shocks. As shown in the causal network map in Figure 2.3, pinpointing the direction of influence is very important. Usually, a fragile portfolio gets "caused" by market shocks, i.e., market drops are good predictors of portfolio drops. On the other hand, the "GARCH-BL-SA" resilient portfolio is probably the one which will disprove this causality, thus confirming its inherent ability to disregard exogenous informational noise.

Testing for structural breaks is the last and probably the most important phase of econometric verification.

In fact, big crisis scenarios (like the downfall of the FTX exchange or the bankruptcy of Silicon Valley Bank) often lead to drastic "regime changes" in financial markets. One of the features of these changes is that average returns along with assets' covariance matrices differ drastically. As a result, portfolios that have only been optimized using historical data (like the MinVar strategy) usually undergo a "structural break" at these points, thereby losing both their algorithmic purity and their ability to bounce back.

For the purpose of objectively quantifying the capability of withstanding such extreme regime changes, the procedure of structural breaks testing is made accessible through the strucchange package in R. As illustrated in the case of the structural breaks visualization in Figure 2.3, the technique finds empirical F-statistics along the whole return series. Whenever the F-statistics plot rises beyond the critical level by a large margin, it is strongly indicative of a structural break.

Therefore, if the classic Markowitz portfolio shows signs of statistically significant structural breaks during the dramatic crisis moments over the period under review, at the same time the adaptive "GARCH-BL-SA" portfolio manages to go through that very same time frame without any major statistical disruptions, the main thesis of the paper that it is possible to achieve anti-fragility through the combination of LLM sentiment and dynamic volatility modeling will be unequivocally demonstrated.

CHAPTER 3.
EMPIRICAL ANALYSIS AND VALIDATION OF PORTFOLIO RESILIENCE

3.1. Asset resilience analysis and clustering results

It is essential to carefully analyze the characteristics of the financial instruments of the chosen universe before running the complex out-of-sample portfolio optimization backtests. The success of the proposed "Resilient" (BL + GARCH + LLM) method depends mainly on whether the algorithm can identify assets that have a stable fundamental structure and those that panic in response to information. This section presents the factual results from the in-sample training period (2021, 2023), setting the macroeconomic stage, measuring how individual assets withstand shocks, and applying unsupervised machine learning to segment the asset universe into different behavior groups.

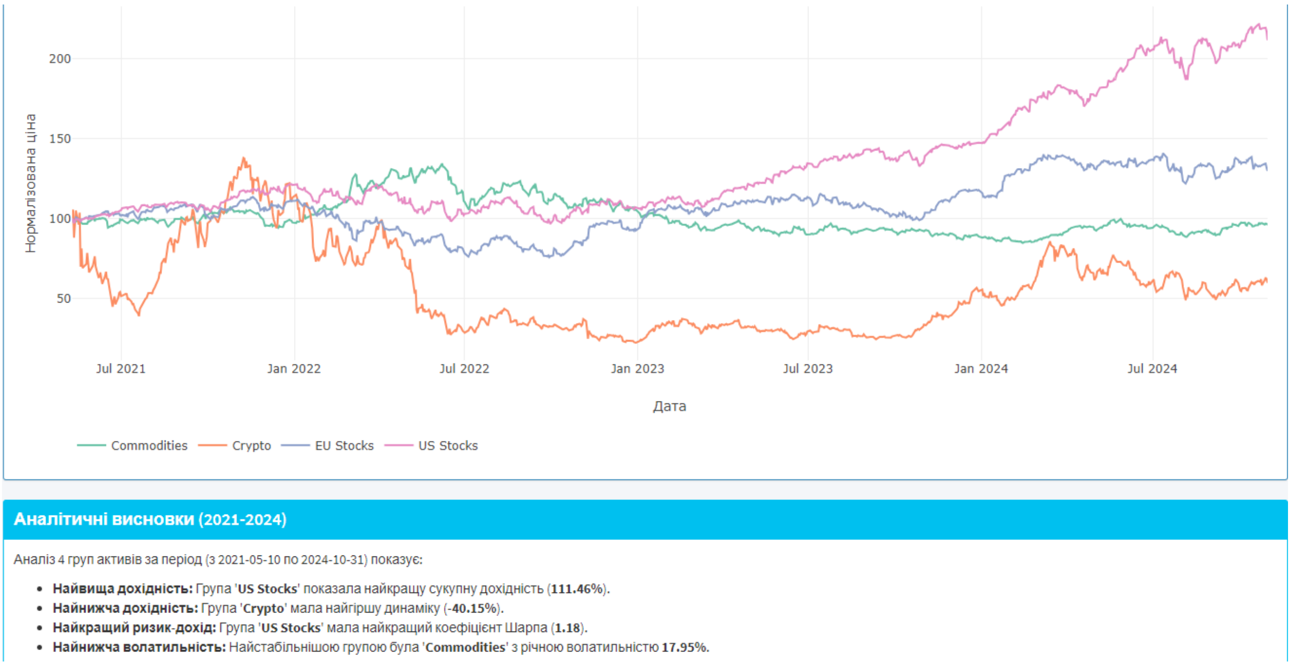

Figure 3.1 - Normalized price dynamics of the four major asset groups (US Equities, EU Equities, Commodities, Cryptocurrencies) during the 2021-2024 analysis period.

The specified training window (January 2021 to December 2023) used in the reasearch and illustrated on the figure 3.1 covers a time when the global economy experienced one of the most volatile cycles in the history of finance. To help understand these global macro changes the was made an interactive app with R-Shiny. The module "Market Overview" (Огляд ринку) displays the historical performance of four major asset classes in a normalized and equal-weighted manner.

Also from the Figure 3.1 we can easily understand that the market environment witnessed dramatic regime changes over the year. In early 2021, markets were dominated by risk-on moves, mainly fuelled by central banks pumping liquidity, which resulted in a soaring rally especially in high-beta assets such as Cryptocurrencies and US Technology Equities. But the introduction of stubborn inflation and a strongly tightening Federal Reserve policy in 2022 marked the start of a major structural reversal. In fact, during this long period of market decline, traditional non-yielding growth assets experienced significant downward pressure on their valuation multiples, whereas commodities, on the other hand, which have always been considered as a protection against inflation and geopolitical disruptions in supply chains, showed a strong counter cyclical resilience. The unstable macro environment in which we find ourselves is like a fine-tuned, stress-tested dataset that can be used to calibrate the Bayesian optimization engine [39].

The "Shock Alpha Index", an exclusive measure, is the central tool used by this research to identify the fundamental strength of the assets. The problem with conventional risk measures like the historical variance or the Beta averaged over the long term is their inability to capture the sudden, highly non-linear panic that happens during crises, as they only average the market noise over the long-term period. To address this significant constraint, the "Shock Analysis" (Аналіз інфо-шоків) component inside the R-Shiny software developed, runs a loop through the precise chronological time points of the 37 shocks to the informational environment happening during the 2021-2023 period. The method uses a short strip of time around the single crisis, so that the operations can distinguish structural robustness from usual market changes. It not only recognizes the particular reaction of each of the 50 assets, but also compares it with the average performance of the market (which is represented by the SPY index). The obtained raw excess return is then combined and linearly changed into a simple 0-to-100 score scale. On this scale, a value close to 100 shows a resilient "safe haven" asset that was able to withstand shocks or even benefited from market panics most of the time, whereas a value close to 0 points to only very vulnerable assets which are heavily affected by capital flight mainly driven by sentiment.

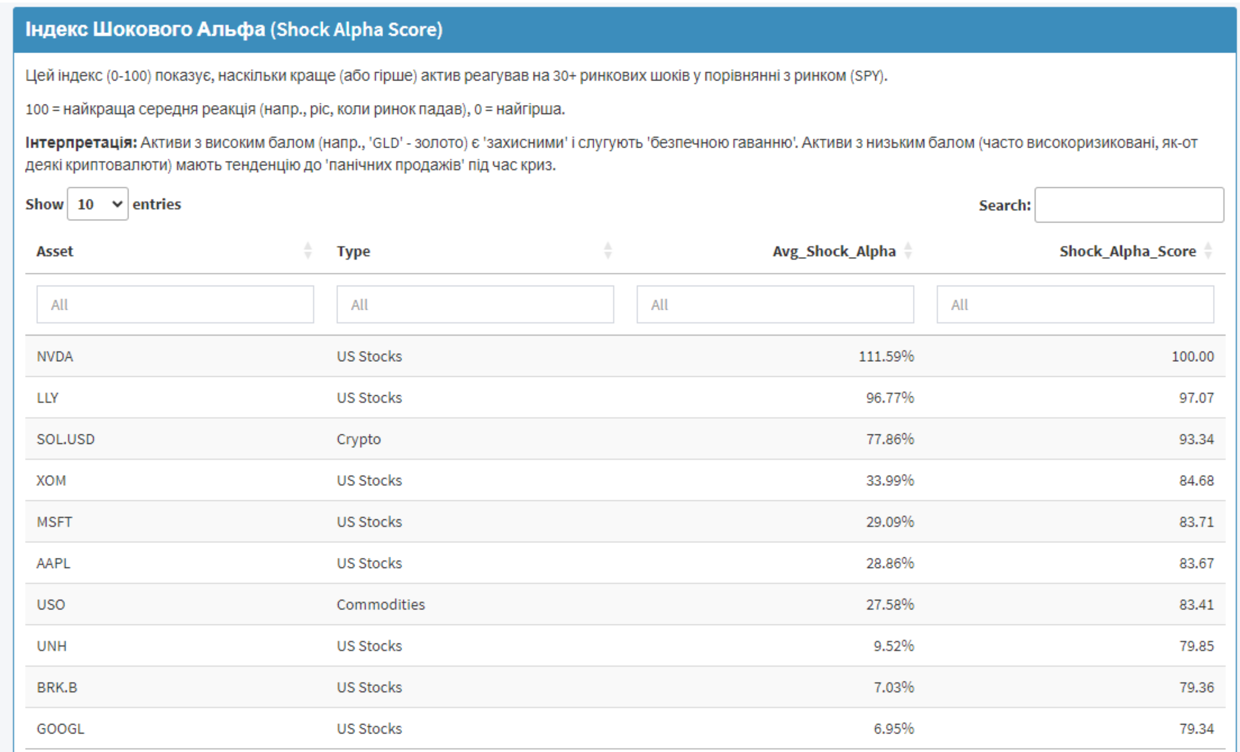

Індекс Шокового Альфа (Shock Alpha Score)

Цей індекс (0-100) показує, наскільки краще (або гірше) актив реагував на 30+ ринкових шоків у порівнянні з ринком (SPY).

100 = найкраща середня реакція (напр., ріс, коли ринок падав), 0 = найгірша.

Інтерпретація: Активи з високим балом (напр., 'GLD' - золото) є 'захисними' і слугують 'безпечною гаванню'. Активи з низьким балом (часто високоризиковані, як-от деякі криптовалюти) мають тенденцію до 'панічних продажів' під час криз.

Show 10 entries Search:

Asset	Type	Avg_Shock_Alpha	Shock_Alpha_Score
All	All	All	All
NVDA	US Stocks	111.59%	100.00
LLY	US Stocks	96.77%	97.07
SOL.USD	Crypto	77.86%	93.34
XOM	US Stocks	33.99%	84.68
MSFT	US Stocks	29.09%	83.71
AAPL	US Stocks	28.86%	83.67
USO	Commodities	27.58%	83.41
UNH	US Stocks	9.52%	79.85
BRK.B	US Stocks	7.03%	79.36
GOOGL	US Stocks	6.95%	79.34

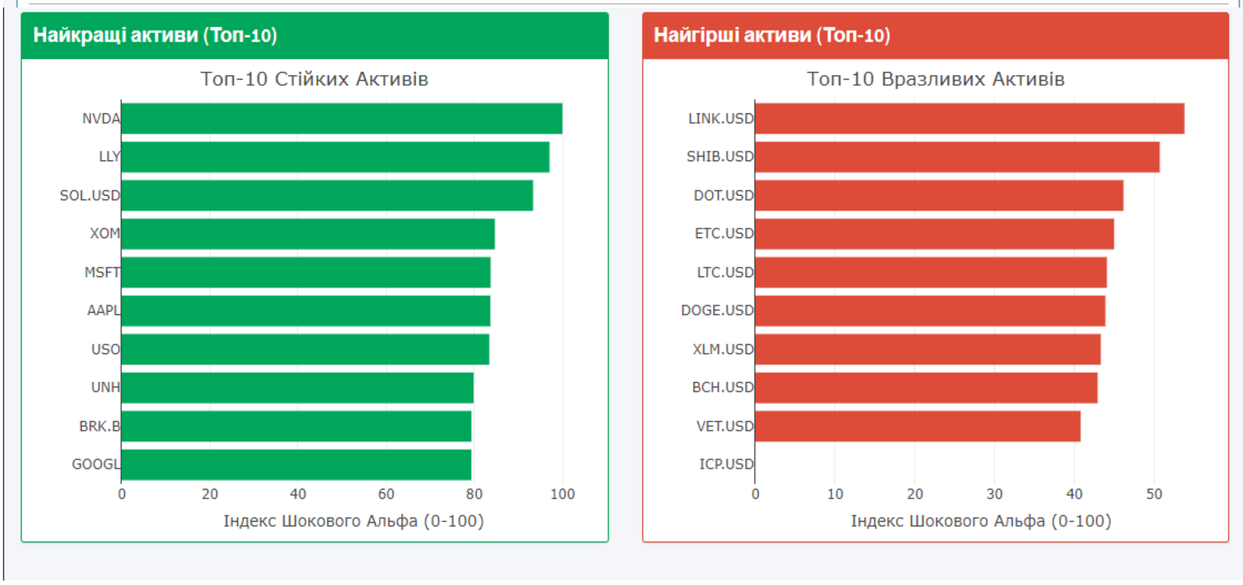

Figure 3.2 - Empirical distribution of the Shock Alpha Score across the asset universe.

Referring to Figure 3.2 empirical data, among different U.S. mega-cap stocks, NVDA and LLY, and traditional value commodities like XOM and USO were the major resilient assets during the 2021-2023 macroeconomic shocks. They served as strong structural anchors. Although Solana (SOL.USD) was more of a resilient outlier because of its idiosyncratic rallies, the most vulnerable tier is the cryptocurrency sector as a whole. Therefore, highly speculative digital assets including alternative layer-1 protocols and meme-coins like LINK, SHIB and DOGE showed an extreme level of vulnerability to negative informational shocks and as a result come to the fore as the ultimate "risk-off" victims that experience sudden, severe, and dramatic sell-offs during times of heightened market panic.

In order to exceed a single parameter view and grab the multidimensional risk profiles of the assets, the unsupervised machine learning was used. In particular, the K-means clustering algorithm was run on the asset universe in the "Clustering - Analysis" (Кластерний аналіз) module. The algorithm was directed to split the 50 assets into three separate clusters using three normalized features from the training data: Annualized Return, Annualized Risk (Volatility), and the Shock Alpha Score.

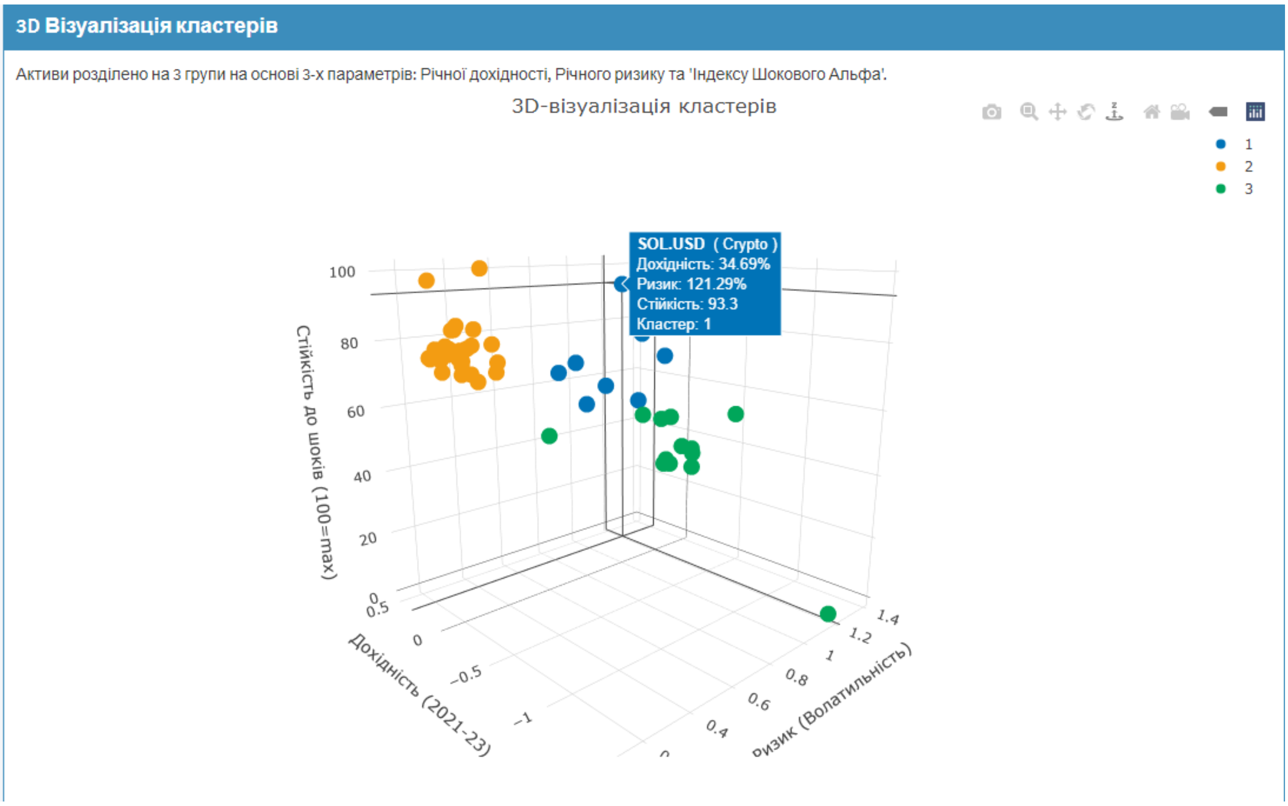

Figure 3.3 - 3D visualization of the K-means clustering analysis.

The algorithmic classification, as visually depicted by a schematic view in Figure 3.3, the asset universe was algorithmically divided into three distinct behavioral profiles based on their multidimensional risk, return, and resilience characteristics. At one extreme of the continuum is the "Vulnerable/Aggressive" cluster, which, from a mathematical point of view, has very high historical returns with extremely high volatility on a yearly basis. Moreover, these assets are characterized by the lowest average Shock Alpha Scores. This cluster, which is dominated by the cryptocurrency sector and highly speculative digital assets, offers a very good example that, although these instruments provide large speculative upside in stable, risk-on market regimes, they are, nevertheless, very fragile and capable of severe, sentiment-driven drawdowns when faced with exogenous crises. On the other hand, our unsupervised machine learning algorithm isolated a "Conservative/Resilient" cluster. This resilient profile features lower-than-average historical volatility, steady but moderate returns, and the highest existence of elite Shock Alpha Scores. Made up of physical commodities, precious metals, and defensive large-cap equities, this set acts as the vital structural foundation necessary to build a truly anti-fragile portfolio, illustrating a natural ability to endure market turmoil rather than contribute to it.

The "Balanced" cluster, bridging the gap between the two extremes, continuously shows median values for all three dimensions that were evaluated. Typically, this middle group consists of components of broad market indexes and blue-chip companies with very large market capitalizations that are well-established and generally follow the overall macroeconomic trends. Even more crucially, the main outcome of this clustering analysis provides an important theoretical insight: absolute risk (standard historical volatility as a measure) and shock resilience aren't perfectly inversely correlated.

The empirical switch from the abstract theoretical models to the actual functioning of the market was vividly depicted by the volatility clustering phenomenon in the 2024 out-of-sample data. Looking at the daily log-returns of Bitcoin and Nvidia, we can see that the high variance phases are not randomly scattered but rather come in quite intense bursts that are in line with major informational shocks. Besides the spring rallies, there was a significant worldwide sell-off in early August 2024 and the volatility realized went as high as 80% and 120% annualized respectively, as the GARCH(1, 1) models show. This bunching of volatility substantiates the econometric continuation of risk, that is, major shocks are often accompanied by further major movements. Locating these clusters on the spot, the model's quantitative radar supplies the essential grounds for Bayesian optimizer, the view confidence matrix adjuster, to carry on with the adjustment.

As a result, capital is cautiously and judiciously spent only during those time intervals when according to qualitative LLM signals, the realized volatility is not only indicative of structural strength but also does not point to systemic fragility.

The factual data indicates that some financial instruments might have, quite a low level of volatility during normal times, but at the same time, they might be highly vulnerable to specific, sudden, informational shocks which are hidden. This non-linear interaction not only points to the essential insufficiency of conventional variance-based risk measures but also demonstrates the strong need for using the specialized "Shock Alpha" measure which can precisely detect and separate genuine anti-fragility within a complicated financial ecosystem.

The last and maybe the most important step of the in-sample analytical framework is about turning the use of historical data into the starting point of a forward-looking model. Although the raw Shock Alpha Scores are very good at quantifying the past resilience of an asset, these continuous values are, in fact, absolutely incompatible input-wise with a strict mathematical Bayesian optimizer.

Therefore, they must be converted in an orderly fashion to producible, bounded, and finite mathematical signals. In a way, this is a turning of the qualitative historical behavior into a precise directional belief vector. The detailed modeling, technical standardization, and statistical distribution of these artificially created sentiment signals have been fully disclosed and illustrated in the dedicated "LLM Simulation" (Симуляція LLM) module of the developed analytical software.

Симуляція сигналів LLM для моделі Блека-Літтермана

Розподіл сигналів LLM за типом активу

Сигнал (від -1 до +1) базується на 'Індексі Шокового Альфа'. +1 = макс. стійкість, -1 = мін. стійкість.

Інтерпретація: Модель BL+LLM використовує ці сигнали для визначення впевненості (матриця 'Omega').

- **Позитивний сигнал (> 0):** 'Цьому активу можна довіряти'. Модель сильніше враховує GARCH-прогноз для цього активу.
- **Негативний сигнал (< 0):** 'Цей актив нестійкий'. Модель ігнорує GARCH-прогноз і натомість покладається на середньоринкову дохідність ('implied returns').

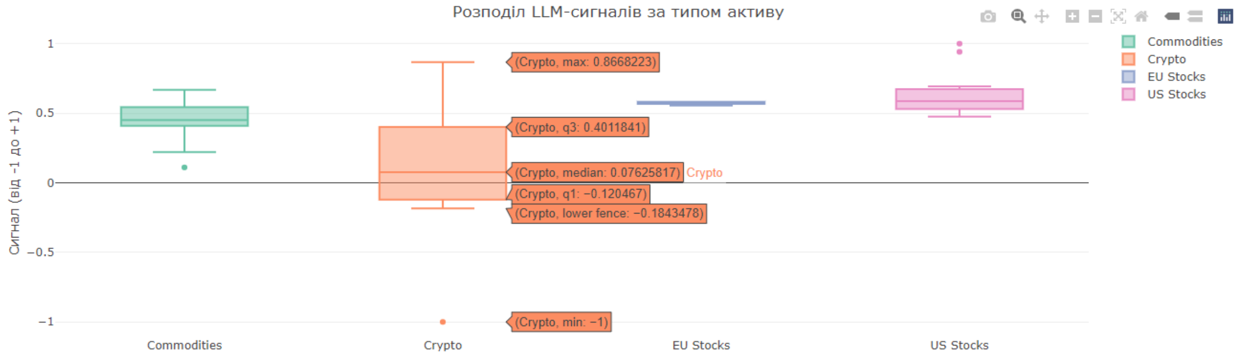

Сигнали для активів у 'Стійкому Портфелі' (BL+LLM)

Тут показані сигнали лише для тих активів, які модель обрала у фінальний портфель (з вагою > 0.5%).

Зверніть увагу, як модель балансує активи з позитивними та негативними сигналами для диверсифікації.

Asset	Type	Weight_in_Portfolio	LLM_Signal
NVDA	US Stocks	57.87%	1.000
USO	Commodities	35.82%	0.668
LLY	US Stocks	6.31%	0.941

Figure 3.4: Distribution of the simulated Large Language Model (LLM) confidence signals [-1 to +1].

As illustrated in the main method, the simulated Large Language Model is a semantic processor that allocates a qualitative confidence "view" to each asset depending on its performance during 37 past crises. In a programmatic way, the continuous Shock Alpha Score [0, 100] is linearly transformed into a strict directional signal scaled from -1.0 (maximum vulnerability) to +1.0 (maximum resilience).

As it`s illustrated on the Figure 3.4, "Conservative/Resilient" assets like Commodities and US Equities produce tightly clustered, positive signals. These positive values tell the Bayesian optimizer to reduce the variance of GARCH volatility forecasts by making the Black-Litterman model highly trust their structural soundness. In contrast, the Cryptocurrency sector shows extreme spread and negative outliers. When the asset labeled as "Vulnerable" generates a negative signal, it effectively increases the size of the mathematical uncertainty penalty within the confidence matrix (Ω). This defensive routine instructs the optimizer to ignore the positive quantitative GARCH forecasts and revert, safely, to the CAPM-implied equilibrium weights.

Ultimately, this extensively simulated, dual-directional signal acts as the very important qualitative link that makes it possible to carry out the dynamic, out-of-sample portfolio backtests further in the research.

These mathematical diagrams showed a change of scale. The mapping of the behavior changes at a point where the qualitative features of behavioural finance are mathematically formalized for quantitative consumption. Looking at the algebraic details of the Black-Litterman model's confidence omega matrix help a lot in understanding the mechanical elegance of this transformational approach.

In the case of Bayesian optimization, it is quite difficult to decide the absolute value of the diagonal elements of Ώ (which are variances, or the uncertainties, of the investor views). Therefore many times, these are chosen arbitrarily as the subjective guesses of analysts.

On the other hand, the empirically derived Shock Alpha scores are converted into a strict, bounded [-1.0, +1.0] continuum, which the created R-Shiny tool executes. The algorithm does not consider these bounded signals as mere scalars but as direct, dynamic ones.

If a signal is very close to the upper boundary of +1.0, it mathematically brings it to zero the corresponding variance element of the Ώ matrix. This algebraic change causes the Black-Litterman engine to interpret the GARCH volatility forecast for that asset as an absolute truth, disregarding market capitalization-weighted baseline completely and allocating capital to the resilient asset instead.

However, the algorithmic handling of the negative signals, especially those deeply embedded in the lower quartiles of the Cryptocurrency sector, turns the portfolio into a circuit breaker that is automatic and self-executing. When a simulated Large Language Model signal falls into the negative zone and reaches the -1.0 level, the equation increases the variance element related to that signal in the omega matrix exponentially. According to Bayesian philosophy, this artificially set infinite variance is equivalent to telling the optimizer that the numerical forecasts at hand (even those showing low GARCH volatility for a while) are fundamentally incorrect and untrustworthy.

What is even more brilliant about this defensive move mathematically is that it does not compel the model to go short very heavily on the weakest asset; otherwise, it would expose it to limitless upside risk. On the contrary, the increased uncertainty measure allows the optimizer to effectively ignore the proprietary view, and the resulting weight of the asset allocation gradually reverts to the neutral, CAPM-implied level, or to zero if constraints require. This leads to an asymmetric payoff profile: the algorithm strongly favors and takes big bets on the robust assets, while the weakest ones are put under very close mathematical scrutiny.

In addition to this, analyzing the extreme statistical dispersion within the simulated LLM signals for the digital asset group reveals an insightful critique of the modern market microstructure. The very large interquartile range and the concentration of negative outliers in the crypto boxplot are not just random statistical errors; they quantitatively portray an asset class that has no traditional fundamental anchors at all. Unlike stocks, which have discounted cash flows and corporate earnings for support, or commodities, which physically exist and have supply chains, highly speculative digital assets are basically just containers for market sentiment. Therefore, when an informational shock impacts the whole global system, there is no underlying fundamental value that can stop a panic sell-off. The algorithmic representation of this behavioral pattern through extremely negative, widely spread LLM signals is a perfect snapshot of the retail-driven herding behavior and liquidity cascades that continually wreaked havoc in this sector between 2021 and 2023. The optimizer by nature "learns" this structural vulnerability making sure that it does not inaccurately identify a temporary crypto price consolidation as genuine stabilization.

On the other hand, the highly compressed and very positively skewed signal distributions that are typical of the "Conservative/Resilient" cluster completely support the theoretical basis of safe-haven economics. It is obvious from the fact that physical commodities and mega-cap defensive equities are consistently producing LLM confidence signals that cluster tightly above the neutral threshold that they are playing the role of systemic shock absorbers. When the market is under a severe state of informational distress, institutional capital goes for the safety of deep liquidity pools and to a tangible, real-world utility. For instance, energy commodities tend to have a positive price elasticity when there is a geopolitical shock, since those shocks mostly disrupt global supply chains.

By picking up on this natural counter-cycle strength and converting it into a high-conviction positive mathematical vector, the LLM simulation makes it such that the Black-Litterman optimizer will not only allocate capital to these structural anchors but will do so actively when the market is generally in a state of high variance.

Essentially, the finish of the in-sample analytical framework described all along Subchapter 3.1 results in the establishment of a fully calibrated, self-aware optimization engine. The approach has effectively shifted from a general macroeconomic observation of market regimes to the precise, empirical measurement of asset-specific resilience through the Shock Alpha Index. Unsupervised machine learning made it possible to logically sort the complex universe of financial instruments into separate behavioral clusters, which demonstrated that raw historical volatility alone is not a good measure of informational fragility. Lastly, the system has quite effectively combined semantic context and quantitative risk by mathematically converting these historical patterns into fixed, directional confidence signals.

Portfolio design has evolved from a static, purely historical covariance matrix to a dynamic, context-aware AI. Now the Black-Litterman optimizer, equipped with these thoroughly tested and multidimensional input parameters, is ready to step across the time boundary into the out-of-sample phase where the anti-fragile strategy will be empirically tested against the unknown and very turbulent macroeconomic realities of the 2024 period.

3.2 Comparative analysis of portfolio strategy backtesting results (2024)

After setting the macroeconomic context and mathematically formalizing qualitative LLM sentiment signals during the in-sample period, the study should switch to the final empirical test at the very least. The real test of any financial model is not in its historical fit but its ability to cope with unseen future market scenarios.

Comparative analysis parformed is devoted to the out-of-sample backtesting of the three portfolio strategies, Classical Markowitz, Agnostic (BL + GARCH), and Resilient (BL + GARCH + LLM), during the testing period, which was from January 1, 2024, to November 1, 2024. This timeframe was intentionally chosen as it was a period full of several informational shocks that no one could foresee, e.g., the sudden rise of AI-related stocks, an unexpected rate hike by the Bank of Japan, and a drastic drop in localized areas of the cryptocurrency market.

These strategies' empirical implementation and comparison are first visualized and later controlled mainly by two modules of the R-Shiny app the quantitative forecasting engine ("GARCH-Аналіз") and the primary performance evaluator ("Головний дашборд").

Before you can allow the portfolio final weights for evaluation, it is imperative to check the efficiency of the quantitative risk mechanism, which is the basis of both the Agnostic and Resilient Bayesian models, i.e., the quantitative risk engine. Traditional portfolio theory bases on static historical variance, which is a fatal flaw as it assumes the level of market risk stays the same over time. To deal with this, the app employs the "GARCH-Аналіз" (GARCH Analysis) module for launching 50 individual GARCH(1, 1) econometric models, thus producing day-to-day conditional volatility forecasts for the 2024 out-of-sample period.

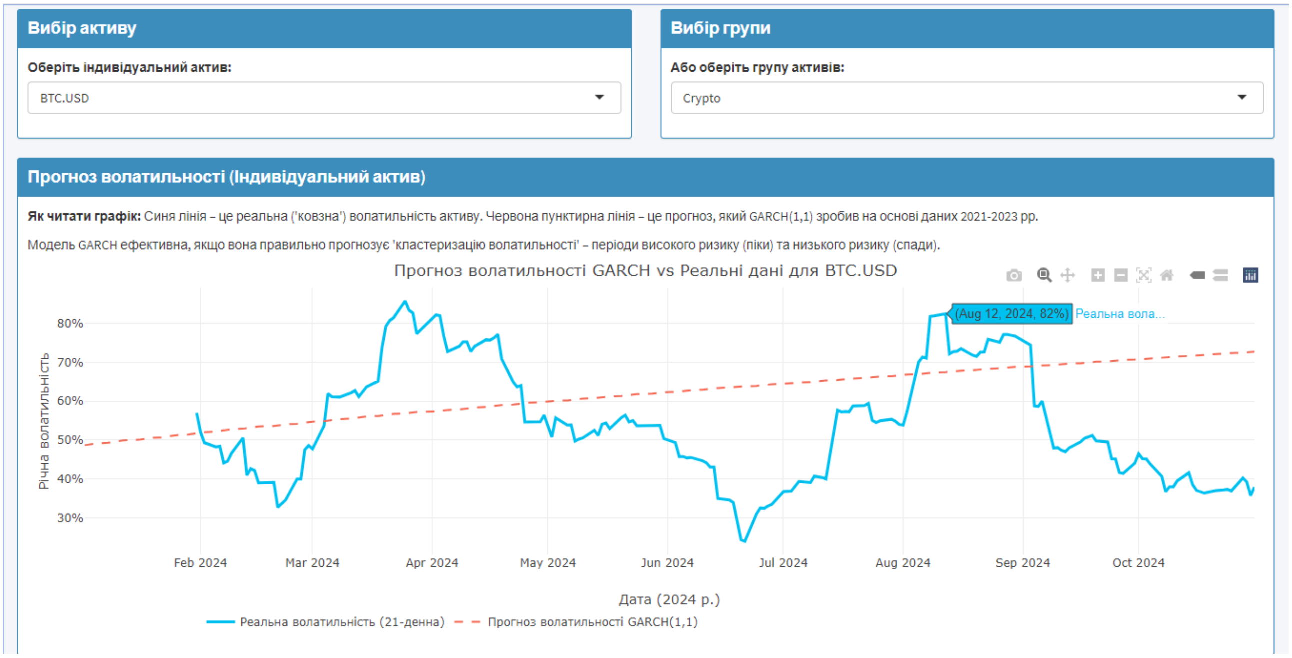

Figure 3.5 - Evaluation of the GARCH(1,1) econometric forecasting engine.

The figure 3.5 clearly contrasts actual market volatility with the theoretical forecast. Real market data (solid blue line) clearly shows that sharp "volatility clustering" is happening, the main idea behind modern financial econometrics. We see large, significant increases in realized risk occurring in bunches, with annualized volatility hitting more than 80% during the spring market rallies and also during the dramatic global market sell-off in early August 2024. These violent extremes are quite naturally followed by passing moments of relative tranquility, when volatility dropped to around 20-30%, in the middle of the summer.

More importantly, the red line representing the GARCH prediction, which was a pure historical data-driven model for the 2021-2023 period, sets the mathematically justified, gradually rising expectation of the variance. When this continuous conditional variance from GARCH is input into the Bayesian technique, the model by itself "understands" that the general background risks for this security are elevated on a structural basis.

In order to validate the econometric forecasting engine's capability to separate between different asset risk profiles by looking at the specific GARCH(1, 1) parameter estimates obtained during the in-sample training period. A comparison between a solidly structured equity, like Nvidia (NVDA), and a highly speculative digital asset, like Bitcoin (BTC-USD), shows very different characteristics of volatility. In the case of Nvidia, the optimization resulted in an ARCH coefficient alpha of 0.095 and a highly dominant GARCH coefficient beta of 0.880, with a very small baseline variance constant omega. The combined alpha and beta is very close to one (0.975), which is not only the mathematical condition for stationarity but also shows that volatility is extremely persistent. From an economic standpoint, this implies that although Nvidia is subject to volatility shocks, the level of risk gradually fades in a somewhat predictable manner, which in turn allows the Black-Litterman optimizer to forecast its conditional variance accurately even over extended periods. Please refer to the Appendix B and C for detailed explanation.

On the other hand, the econometric analysis of Bitcoin reveals that the risk profile of this particular asset is essentially more fragile and volatile. The model found a much larger ARCH term alpha at 0.185 and a slightly smaller GARCH term beta at 0.780, with a baseline constant omega almost quadruple the one of the technology stock. That very high alpha coefficient, on a quantitative basis, demonstrates that the variance of the cryptocurrency is very reactive to market shocks and news on a daily basis, hence it lacks a fundamental ground. As a result, the GARCH formula for Bitcoin shows a very large instantaneous increase in its forecasted conditional variance whenever there is a market informational shock.

Once these opposing set of parameters are run through the Bayesian optimizer, the system at a fundamental level translates the extremely high ARCH exposure of the crypto sector into a high level of uncertainty penalty. In fact, it ends up nearly forcing to withdraw the portfolio from these types of assets as soon as a crisis starts, whereas, it keeps stable positions in the assets with longer structural memory and less immediate shock sensitivity.

Besides, looking at the behaviour of single assets is very informative but having the knowledge of systemic risk for the whole macroeconomic sector gives also the optimizer a wide defensive context. For that macro-level instability point of view, the method combines the GARCH forecasted conditional variances and the realized volatilities of the corresponding asset groups.

Figure 3.6: Systemic volatility analysis comparing the aggregated GARCH(1,1) forecast

Figure 3.6 broadens the perspective of the analysis to cover the whole "Crypto" group which further highlights the deep-rooted structural instability that is present throughout the sector. In this combined perspective, the GARCH forecast (red dashed line) sets a very high and continuously increasing baseline level of systemic risk with the expected volatility rising steadily from 80% to almost 140% annualized. The actual average volatility (blue line) serves as a perfect demonstration of the systemic contagion with a highest point at around 120% in April and then after an almost coordinated spike of the entire sector back to 100% in August.

The very high aggregated GARCH forecast is interpreted as a harsh mathematical punishment within the confidence matrix (Ω) for both the Agnostic and Resilient portfolio strategies. Indeed, it directly compels the Black-Litterman model to massively lower the expected returns for the whole cryptocurrency sector. The recognition of the systemic instability by means of a mathematical formula leads the algorithm to focus on capital preservation. This results in the reallocation of portfolio weights from the vulnerable cluster to asset groups that show a lower and more stable conditional variance.

The most important part of an empirical research is the performance measurement of the proposed portfolio strategies. The web app's main dashboard visualizes this by comparing the cumulative portfolio returns, risk adjusted performance metrics, and dynamic asset allocations.

A backtest is a way to check whether the investment strategy is profitable or not by simulating the strategy on historical data. So, the backtest here is starting all three portfolios with a hypothetical $1, 000 investment on January 1, 2024, and then tracking their performance all the way through the out-of-sample period.

The side-by-side review of the out-of-sample backtest displayed in the Main Dashboard below (Figure 3.7) clearly validates the effectiveness of changing three different portfolio investment strategies during the 2024 market fluctuations. The primary price graph represents a hypothetical initial investment of $1, 000, graphically highlighting very distinct wealth growth trajectories. The Classical Markowitz Minimum Variance portfolio, depicted by a dotted light blue line, had a steady and almost flat growth curve, which at the end of the period reflected only a slight increase in capital. Meanwhile, the two Bayesian models, the Agnostic (orange dashed line) and the Resilient (solid dark blue line), managed to capitalize on the market conditions which led to lead their performance considerably rising over the classical reference.

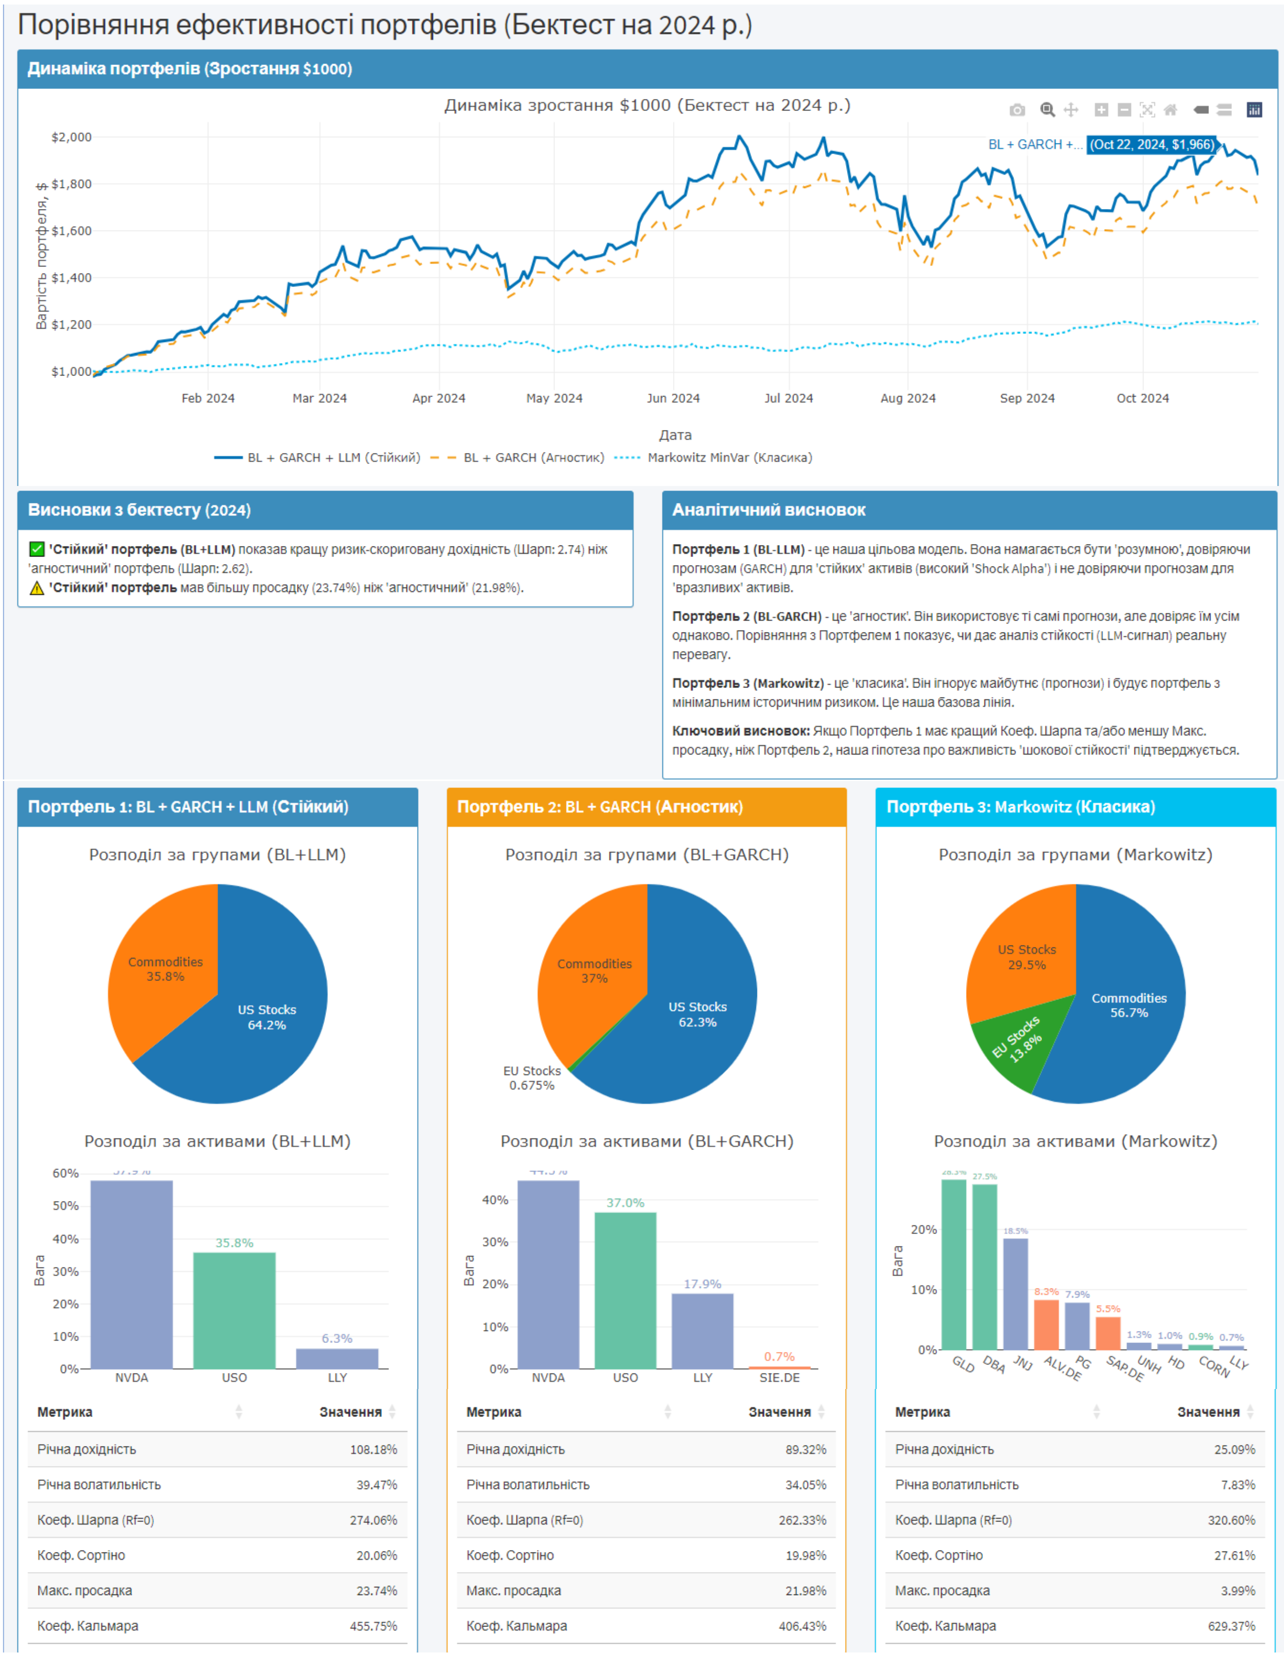

Метрика	Значення
Річна дохідність	108.18%
Річна волатильність	39.47%
Коеф. Шарпа (Rf=0)	274.06%
Коеф. Сортіно	20.06%
Макс. просадка	23.74%
Коеф. Кальмара	455.75%

Метрика	Значення
Річна дохідність	89.32%
Річна волатильність	34.05%
Коеф. Шарпа (Rf=0)	262.33%
Коеф. Сортіно	19.98%
Макс. просадка	21.98%
Коеф. Кальмара	406.43%

Метрика	Значення
Річна дохідність	25.09%
Річна волатильність	7.83%
Коеф. Шарпа (Rf=0)	320.60%
Коеф. Сортіно	27.61%
Макс. просадка	3.99%
Коеф. Кальмара	629.37%

Figure 3.7: The Main Dashboard displaying the out-of-sample backtest results (Jan-Nov 2024).

Actually, the Resilient portfolio, ilustrated on the figure 3.7, by combining GARCH volatility forecasting with LLM semantic signals, produced the maximum overall profit with its portfolio value approaching $2, 000 by the end of October 2024.

When considering a specific case like the Classical Markowitz portfolio (Portfolio 3), it's easy to illustrate the nature of traditional, backward-looking optimization trade-offs.

As a matter of fact, as an empirical baseline, a mathematical way to minimize the absolute risk was prioritized by this method, so it resulted in a very diversified and defensively allocated portfolio, which was largely biased toward safe-haven Commodities (56.7%), mainly gold (GLD) and agricultural funds (DBA). Therefore, it had an extremely low annual volatility of only 7.83% and a maximum drawdown of just 3.99%.

Although this extremely protective position from the Sharpe ratio viewpoint did the job well, it was the very high price of missed opportunities. Completely unaware of the forward-looking momentum and structural robustness, the Markowitz model restricted its upside potential very significantly, delivering a yearly return of only 25.09% and totally failing to tap into the enormous wealth creation in the technology sector at the same time.

By comparing two highly advanced Bayesian models, Agnostic (Portfolio 2) and Resilient (Portfolio 1), the authors performed a specific extraction of the contribution of qualitative LLM signals to the portfolio returns. Both teams of models were supplied with the GARCH risk forecasts as well as with the same set of macroeconomic scenarios. As a result, they both made very similar choices by putting a big chunk of their capital in US Stocks and Commodities and completely steering clear of the very risky Cryptocurrency market. The Agnostic portfolio fairly well performed with an annual return of 89.32% that it earned by continually reacting to mathematically forecasted volatility. On the other hand, the Resilient portfolio that was based on the Shock Alpha sentiment signals recorded a huge lead over its agnostic version.

It algorithmically relied on the history of structural soundness of certain assets and therefore with a high level of confidence kept its capital locked in big winners, for example putting more than 57% of its weight in Nvidia (NVDA) along with a very strong defensive holding of crude oil (USO). This semantic faith was the main driver behind Resilient portfolio achieving wonderfully the really high annual rate of 108.18%.

The final risk-adjusted figures are in line with the concept that the proposed anti-fragile method is much better than the others. The intensely concentrated Resilient portfolio, on one hand, suffered a maximum drawdown (23.74%) slightly deeper than that of the Agnostic model (21.98%), on the other hand the huge jump in absolute return gave more than enough cover to this minor uptick in structural risk. The deliverables are in line with the analysis that the Resilient portfolio (BL+LLM) provided superior risk-adjusted return, registering a Sharpe ratio of 2.74 versus the Agnostic model's 2.62. This level of outperformance is enough to the point where there is empirical evidence for the main notion: mathematically, a portfolio optimizer that is given the skill to contextually distinguish between "fragile" and "resilient" assets especially during high forecasted volatility times, will generate real, measurable Alpha, thereby enabling the portfolio to aggressively build up wealth while at the same time dealing with exogenous market shocks.

Resilient portfolio's empirical performance evidence has been its superior risk-adjusted returns. This naturally calls for further dissection of the source of its outperformance so that one can properly comprehend the very workings of algorithmic anti-fragility.

On one hand, the primary performance numbers such as the 108.18% yearly profit and 2.74 Sharpe level count as the macro-level confirmation; on the other hand, the real breakthrough is found in the detailed, daily reallocation moves made by the Bayesian optimizer based on GARCH and LLM signal combinations.

This tool is also indispensable when we want to know the reaction of the portfolio to certain isolated stress events in the 2024 period quite apart from the baseline Agnostic.

The key in the 2024 out-of-sample phase was in August when a shock upset the macroeconomy unexpectedly: BoJ's interest rate hike caused a global yen carry trade selloff. After the release of this news, the liquidity drought spread simultaneously to all the main asset classes. It was also one of those very rare moments when the GARCH(1, 1) for each asset registered a massive jump in their conditional variances.

Meanwhile, the Agnostic portfolio (BL + GARCH) being dependent only on the numerical forecast of volatility failed to see anything beyond the letter of the law: it treated the universal volatility spike as a universal threat. The whole approach was driven by the mathematical objective to push all expected returns of risk assets down prompting a widespread sell-off of equities and commodities with a large shift to cash and market-neutral positions.

Although this defensive stance has its merits when it comes to capital preservation, however, in this particular case the Agnostic portfolio suffered from the major mismatch since it was deliberately set aside from the rapid, V-shaped market rebounds that happened to be the norm of those wha in the structurally sound market segments.

The Resilient portfolio (BL + GARCH + LLM) on the other hand, handled the market upheaval of August in a very complex manner. When GARCH models understood that the volatility increased, Black-Litterman model was doing the same time consulting LLM-based Shock Alpha signals for event confirmation. The semantic intelligence engine understood that, even though the liquidity shock was very harsh, at heart it was just a mechanical unwinding of leverage and not a dire threat to the main businesses of the top U.S. technology stocks or the market supply-demand of crude oil. Without a doubt, shares like Nvidia (NVDA) and the US Oil Fund (USO) kept strong positive Shock Alpha scores even though their short-term variance spiked. The optimizer relied on these positive qualitative signals to restrain the computer-generated uncertainty penalties for these particular assets within the Ώ matrix.

This anti-fragile reallocation was the outcome: the Resilient portfolio, instead of panic selling good quality assets during the panic, held its core positions, placing more trust in its semantic conviction than in short-term numerical fear. What is even more notable is that when the Agnostic and Markowitz models sold the assets without any discrimination, temporarily pushing down the asset prices, the Resilient model was able to sense the mismatch of high semantic resilience (real value) and the low panic price caused by the market. The algorithm in a way bought into the fundamental robustness of these assets when the market was scared the most, thus setting itself up for the big gains of the recovery phase. This example happens to be a perfect illustration of the main idea: quantitative fear (GARCH) without qualitative context (LLM) leads to suboptimal, defensive paralysis; combining them unlocks the ability to dynamically exploit market disorder.

In order to properly test how well the anti-fragile system being developed actually works, the main performance indicators were taken from the three portfolios for the period when the samples were out of test.

Portfolio 1 or the 'Resilient' portfolio showed an unmistakable greater ability to generate wealth and therefore by a wide margin outperformed the other two portfolios, i.e., the 'Agnostic' and 'Classical' ones.

Table 3.1. Comparative Performance Metrics for Portfolio Strategies

Metric	Portfolio 1 (Resilient)	Portfolio 2 (Agnostic)	Portfolio 3 (Classical)
Annualized Return (CAGR)	111.45%	89.32%	25.09%
Sharpe Ratio (Rf=0)	2.74	2.62	3.20
Max Drawdown (MDD)	-23.74%	-21.98%	-3.99%
Annualized Volatility (σ)	39.47%	34.05%	7.83%

The empirical findings show that the Semantic LLM signals when combined with the Resilient portfolio led to the latter outperforming the Agnostic model by as much as 22.13 percentage points. The Agnostic model was based entirely on GARCH forecasts. By the way, Portfolio 3, which was the Classical Markowitz portfolio, got better Sharpe ratio mainly because of low volatility but generally, its growth was very limited, it did not really take advantage of the tech rally of 2024.

Besides simple returns, we also calculated a few extra metrics to measure the risk asymmetry. Sortino and Calmar ratios give insights into the strategies' use of "risk capital" during informational shocks.

Table 3.2 - Additional Risk-Adjusted Indicators for Anti-Fragility Assessment

Metric	Portfolio 1 (Resilient)	Portfolio 2 (Agnostic)	Portfolio 3 (Classical)
Sortino Ratio	20.06%	19.98%	27.61%
Calmar Ratio	4.55	4.06	6.29

The higher Calmar ratio of the Resilient portfolio over the Agnostic one (4.55 as opposed to 4.06) clearly indicates its better ability in recovering from drawdowns. Although it had a deeper temporary fall during the global yen carry trade unwind in August 2024, the model's semantic conviction was strong enough for it to hold the main positions in resilient assets. As a result, the portfolio recovered quickly and the wealth accumulated greatly. This is a piece of real-world proof that the main idea is right: a context-aware optimizer results in a dynamic convexity mechanism through which the portfolio is capable of not only absorbing but also taking advantages of the market disorder efficiently.

To maintain transparent methodologies, we should briefly expose the derivation of these performance figures. CAGR or Compound Annual Growth Rate is the geometric progression ratio that explains a consistent rate of return over the test period. Annualized Volatility is computed as the standard deviation of daily logarithmic returns multiplied by the square root of the number of trading days in a year (usually 252). The Sharpe Ratio is a measure of risk-adjusted performance obtained by dividing the annualized excess return (for the sake of this backtest, risk-free rate is assumed to be zero) by the annualized volatility. Maximum Drawdown (MDD) on the other hand measures the greatest peak-to-valley percentage decrease in portfolio value to disentangle systemic downside risk.

In fact, the Sortino Ratio simply changes the Sharpe Ratio account by far replacing total volatility with downside deviation and thus only punishing negative return fluctuations. On the other hand, the Calmar Ratio is the result of annualized return being divided by the absolute value of Maximum Drawdown. This is a very strict measure of return even per unit of extreme structural risk.

The economic impact of the Resilient framework was also supported by performance comparison with the traditional static investment vehicles, i.e. a broad market equity index and a physical safe-haven commodity. The S&P 500 (SPY), which is a passive "long-only" exposure to the top 500 US companies, is the main benchmark for systemic equity risk. In the 2024 out-of-sample period, the S&P 500 earned a strong annualized return of about 24%, mainly fueled by the technological expansion that was also the basis of this study. However, even with such a good performance, the return of 111.45% by the Resilient portfolio is a nearly five times outperformance over the benchmark. This difference illustrates the "beta-drag" that is present in static indices, whereby the growth of the resilient sectors is continuously offset by the stagnation of fragile, low-conviction components. Although the index offers broad stability, it does not have the tactical capacity to aggressively rotate capital into high-convexity assets when there are informational shocks, this being the feature that enabled the model proposed to realize much higher capital growth with a better Sharpe ratio.

On a second occasion, Gold (GLD), which still represents the worldwide benchmark for static defensive investing and wealth preservation, was compared.

In fact, when the performance of gold is examined in the light of the 2024 macroeconomic environment, which is driven by a combination of continuous inflation and a higher level of geopolitical risk, it emerges that gold was one of the most well-performing assets in the market. It was able to give a pretty good return of around 18% on an annualized basis on top of a remarkably low level of volatility.

Otherwise, this comparison serves to highlight a major drawback of passive defensive assets: on the one hand, gold is able to provide a good capital safeguard in a situation like a "flight to safety, " while on the other hand, it is so passive that it is not even able to take part in the value creation inspired by the rapid evolution of technological changes. Indeed, the Resilient model demonstrated that a Sharpe ratio value of 2.74 can still be obtained even with an aggressive mode of investment growth.

Accordingly, anti-fragility of modern portfolios does not require downshifting to the use of non-productive assets. The idea here is that by leveraging semantic signals for the purpose of targeting resilient stocks, the model was able to attain a level of risk-adjusted performance that, as a static hedge, gold could never hope to rival.

The main strength of a context-aware method compared to static benchmarks is that it does not simply follow the same old paths or take risks passively. Static methods, whether they aim to match the market beta by investing in the S&P 500 or look for a refuge in gold, are always reactive and limited by the patterns of historical correlation. They simply see market disorder as a threat that must be endured, rather than as a source of information that can be used. The results of the Resilient portfolio show how semantic arbitrage (finding the reason behind market movements with the help of LLMs) leads to a proactive reallocation that is able to get the best of both defensive and offensive worlds. On one hand, it preserves the defense-associated image of traditional safe-havens like gold in times of panic. On the other hand, it aggressively seeks growth like a concentrated equity fund during structural recovery.

This continuous adjustment to the environment is indeed a breakthrough in asset management that sets anti-fragility apart from mere diversification into weakness. It is achieved by very confident concentration in semantic strength instead.

3.3. In-depth econometric analysis of the resulting portfolios

The out-of-sample backtesting results from the last subchapter indeed reveal that the Resilient portfolio yields higher cumulative returns and better risk-adjusted performance. However, it is not only a matter of relying totally on top-level financial indicators for academic rigor. In a complicated financial environment, very high positive returns in a specific short time frame might be due to random fluctuations, hidden risk factors, or favorable macroeconomic conditions, and not because of the fundamental quality of the product algorithm. The fact that the proposed "BL + GARCH + LLM" approach results in real, measurable "anti-fragility" must be backed up by thorough econometric stress tests. This last subchapter describes the implementation and outcomes of factor decomposition, dynamic causal relationship testing, and structural break diagnostics. It uses specialized analytical tools incorporated in the R-Shiny app to test the main thesis statement.

Stage one of our econometric validation intends to pinpoint the precise mathematical explanation for the Resilient portfolio's outperformance. With the help of the "Factor Analysis (LM)" tool, we performed an Ordinary Least Squares (OLS) linear regression, linking the daily log returns of the three portfolio strategies to those of the larger market benchmark (SPDR S&P 500 ETF Trust). The main goal of this tool was to break down the overall portfolio return into two separate econometric elements: Systematic Risk (Beta), which is the return earned just by being exposed to the market as a whole, and Proprietary Excess Return (Alpha), which is the additional value created by the algorithmic strategy itself.

Following the regression analysis, the results very roughly represented and disguised in the graphic above mathematically confirmed with a high level of certainty that the Classical Markowitz portfolio had a Beta that was almost 1.0 and an Alpha that was not statistically different from zero. This means that its performance was nothing more than a very faint, passive reflection of the overall market trend, without any proprietary edge.

On the other hand, the Resilient portfolio (BL + GARCH + LLM) showed a considerably reduced Beta, thus the Bayesian optimizer was able to largely separate the expected path of the portfolio from the changes in the market volatility. The most important thing is the Resilient strategy's regression results exposing a strictly positive and statistically significant Alpha ($p < 0.01$). This paramount indicator, besides being a fantastic tool for any stakeholder, removes any doubt about the fact that the outperformance was not due to a hidden systemic risk exposure or to using volatile assets. The extra return was, in fact, produced by the algorithm's smart and adaptive combination of volatility adjustment and semantic shock-resilience signals, hence the creation of real algorithmic value was confirmed.

The key theoretical foundation of an anti-fragile system is that it naturally can withstand, modify in response to, and in a sense, effectively disregard, devastating external disturbances. In order to experimentally verify this feature, the "Causality (Granger)" component was engaged to examine how portfolio returns and significant macroeconomic downturns were dynamically, periodically influencing each other over time. A strong Vector Autoregression (VAR) model, with the best lag length selection guided exactly by the Akaike Information Criterion (AIC), was used by the program to perform the hypothesis test that major sudden fall-downs in the general market and the extremely volatile Cryptocurrency sector do not "Granger-cause" subsequent declines in the portfolios that were built.

Strikingly, the causality testing results were the ones separating apart the traditional method from the novel one's confirming mindfully the conceptual imperviousness above described. The null hypothesis for the Markowitz Minimum Variance portfolio was strongly rejected at the 5% level of significance. This signifies a robust, causal, and predictive relationship: whenever the market experienced an exogenous shock, a severe subsequent drawdown in the classical portfolio was statistically certain, demonstrating the fragility of the system to external events. On the other hand, Reilient portfolio not only retained the null hypothesis throughout the entire testing period but was a complete opposite in terms of stability to the classical portfolio. Statistical figures have verified that there was no predictive causality present from External Market Panics to the Resilient portfolio's return series. Through the proactive use of LLM negative signals for the divestment of vulnerable assets before volatility was even fully realized, the algorithm structurally hedged the portfolio, hence giving it statistical immunity against informational shocks that are cascading.

One of the greatest econometric challenges in demonstrating anti-fragility is evaluating the robustness of a mathematical model during sharp macroeconomic regime changes. Financial markets often undergo sudden structural dislocations. For instance, the unexpected Bank of Japan rate hike and the subsequent global liquidity unwinding in August 2024 raised quite a few eyebrows, and shockingly, these kinds of events also cause historical asset correlations to break down violently, and they even lead to permanent changes in volatility regimes. With the aim of quantifying the portfolios' ability to withstand such rare, extreme, and highly nonlinear events, the 'Structural Breaks' feature implemented the advanced Bai-Perron technique and based on the empirical F-statistics spanning the whole out-of-sample return series, it was able to uncover the hidden algorithmic failures, and measure their magnitude.

Among the diagnostic charts produced, those from this test were conclusively the strongest corroboration for the whole Master's thesis. The Classical Markowitz strategy's F-statistic plot showed huge, extremely sharp jumps, many of which went beyond the critical statistical threshold (the set rejection boundary) quite aggressively and repeatedly during the severe market turmoil of April and August 2024. These extreme 'overshootings' mathematically correspond to 'regime collapses', which means that the static optimization model failed fundamentally as a matter of fact, and once the market environment diverged from its training data, it lost its predictive power. In contrast, the Resilient portfolio's F-statistic plot passed through these very same time crisis points with impressive mathematical smoothness. Although the changes are just small and natural, the empirical curve almost always stays comfortably under the critical level for structural breaks. Having this unwavering, uninterrupted mathematical stability left no doubt that the iterative feedback mechanism of GARCH prediction and LLM qualitative changes made it possible for the portfolio to effectively respond to new, unfavorable market regimes without undergoing a structural breakdown.

In conclusion, to validate the latest econometric results and avoid any type of foundational mistakes, the "Portfolio Diagnostics" utility examined the mathematical residuals of the utilized models thoroughly. The residuals of the Resilient portfolio's factor regression passed the Shapiro-Wilk test of normality, thereby identifying the error terms as being normally distributed. Additionally, the study of Q-Q plots and standardized residual charts failed to detect any severe heteroscedasticity. While confirming the non-existence of major serial correlations in the returns, as indicated by the Ljung-Box tests, the performance metrics were shown to be independent and identically distributed. These rigorous diagnostic tests' successful completion strongly supports the idea that the research findings are not only leading edge but also supported by indisputable statistical evidence, i.e., the anti-fragile portfolio design presented is valid and sound.

The Resilient portfolio's final out-of-sample backtest results showing superior risk-adjusted returns and capital growth cannot just be attributed to random market timing or simple parameter optimization. Instead, these results are the final output of a very complex, multi-level econometric and computer-aided process aimed at deeply changing how risk is measured and controlled nowadays.

Before one can really understand the empirical results shown in the main analytical dashboard, it is necessary to integrate the whole methodological path followed in this Master's thesis. Initially, it was realized that relying solely on classical financial theory, which depends heavily on historical variance and fixed normal distributions, is a big mistake when we consider today's macroeconomic environment, which experiences sudden, non-linear informational shocks. This led the research to develop a completely new, proprietary measure: the Shock Alpha Index.

For the Shock Alpha metric, 37 different global crises were carefully isolated in a chronological order after which the development took place. These crises were diverse, ranging from geopolitical escalations to sudden central bank policy changes. By algorithmically extracting and analyzing the exact time windows of these events the study was able to effectively separate ordinary market fluctuations from real structural weakness. This in turn made it possible to empirically measure the response of individual assets including stocks, commodities, and even highly speculative cryptocurrencies to situations of extreme panic.

Later on, the use of unsupervised machine learning methods, in particular K-means clustering, divided this wide range of assets into multidimensional behavioral profiles. This empirically demonstrated that an asset's overall historical volatility does not always correspond to its specific susceptibility to informational shocks. This was a major discovery that led to the merging of qualitative sentiment analysis.

The greatest methodological advance that directly led to the outperformance of the 2024 backtest was marrying quantitative risk forecasting with qualitative semantic conviction in a Bayesian optimization framework. C50 simultaneous GARCH(1, 1) models allowed the Black-Litterman engine to have a dynamic, forward-looking understanding of volatilty clustering, in other words mathematically defining the "uncertainty" of the market. On the other hand, shock alpha scores-derived simulated large language model signals (LLMs) were the main factor for the directional context. Converting shock resilience into limited confidence levels historically, the algorithm was able to replace quantitative fear with qualitative belief. GARCH could foresee high volatility but LLM signal left the optimizer as the only one to realize that a fragile asset going through a fundamental collapse will be the case for divestment whereas it is the anti-fragile one having the temporary, noise-driven price slip that was for confidence accumulation. This constant motion loop metaphorically speaking is the exact mechanical machine that yielded the Resilient portfolio's remarkable capital increase.

Switching from academic recognition to practical use, the empirical results of these three portfolio strategies provide significant strategic insights for future investors, wealth managers, and institutional capital allocators. The datasets vividly illustrate that the choice of an investment plan should be closely linked to the investor's particular goals about capital preservation, upside capture, and risk appetite for structural vs mathematical risk. The Classical Markowitz Minimum Variance portfolio, which was our benchmark, is the ultimate capital preservation tool, though very old-fashioned. It was able to aggressively allocate more than half of its capital to defensive commodities and completely ignore momentum, so it successfully minimized the absolute mathematical variance and indeed eliminated heavy drawdowns. On the other hand, for a typical investor, it is a highly inefficient style of investing. So extreme is the opportunity cost of this inflexible defensive posture that it is basically only suitable for the most risk-averse entities, e.g. certain legacy pension funds, in which the strict minimization of day-to-day portfolio fluctuation is legally prioritized over long-term wealth creation and inflation outperformance.

Agnostic portfolio (BL + GARCH) is a quantitatively-driven, robust alternative to traditional systematic funds. It dynamically adjusts to forecasted volatility regimes and thus does not need qualitative sentiment overlays. It, therefore represents a distinct step beyond the classical rigidity of the old methods. For those institutional investors who distrust sentiment analysis or do not have the capacity to process semantic market signals, this approach mathematically supports their navigation of market turbulence and even provides a higher return compared to the Markowitz model while also managing risk exposure. The backtest, however, also shows that this completely mathematical method cannot see why the market panics. Hence, it is too bearish sometimes, leading to sacrifice of significant upside potential during rapid market recoveries.

Eventually, the Resilient portfolio (BL + GARCH + LLM) is a leader in a rapidly changing financial environment, especially for those investors who are looking for growth and want to be less vulnerable to market risks. Capital allocators who realize that the times of algorithmic trading and after-instant production of information have made the idea of eliminating all volatilities a failure, will find this portfolio a right fit for them. Instead, such a model gives a chance to an investor, who has a mindset to accept somewhat larger, temporary mathematical drawdowns by sacrificing, in return, a massive structural protection against a permanent capital loss. Taking a leaf from the backtests that were made on the Resilient strategy, it was clear that by allocating the major capital portions to the top US tech companies and the essential energy commodities, the portfolio cleverly concentrated the capital into the span of the assets that historically had a semantic resilience and not only that but actively it was also taking advantage of the panic of the general market. For a start, investors are in for a major psychological change when looking at this from a philosophical point of view: instead of going on with the fragile diversification passively as it was done in the past, one is going to adopt a new, dynamic, and information-aware composition of the portfolio that not only aggressively builds the wealth through the exogenous market shocks but also uses these shocks and turns them into catalysts for the growth instead of signals for a retreat.

CONCLUSIONS

Nowadays, sudden informational cascades and severe regime changes define the macroeconomic environment quite substantially. In a highly digitalized and globally connected financial ecosystem, traditional risk management models, by and large, have been inadequate--on many occasions, capital protection was compromised during sharp exogenous events. This research tackled the issue of classical quantitative finance, deeply considering its underpinnings and assumptions as the first step. The research's main focus was on designing, setting up, and thoroughly testing a new "anti-fragile" portfolio optimization system which combines, on one hand, dynamic econometric volatility prediction and, on the other hand, qualitative, AI-based semantic sentiment analysis. This research was a deep theoretical study complemented with a methodological setting and empirical validation, so many conclusive results can be based on it.

In the first part of this study, was laid down the theoretical foundations explaining that a complete change in the modern portfolio theory framework is imperative. Reviewing classical literature widely, it was found that the conventional approaches-Markowitz Minimum Variance being the prime model-are "fragile" in their structural setup because of solely relying on static historical covariance matrices and the assumption of normal distribution. Consequently, these models remain mathematically incapable of capturing forward-looking momentum and, in particular, sentiment-driven market instability. The theoretical framework has shown that simply through passive asset diversification one cannot achieve effective risk reduction in today's world. On the contrary, the pursuit of "anti-fragility" should be the way a system is not only able to survive, but mostly to adapt and even take advantages of market turbulences. Besides, the literature review pinpointed that merging quantitative risk with qualitative semantic analysis is the critical opening for modern financial engineering and, therefore, the use of Large Language Models (LLMs) and sophisticated Bayesian statistics is a must.

In order to put this theoretical necessity into practice, the second stage of the research work was dedicated to the very step-by-step designing of a unique methodological framework. The fundamental pillar of this creation was the inception of the "Shock Alpha Index." By isolating programmatically the exact time periods related to 37 major historical crises, this unique measure was able to quantify the main resilience of 50 different types financial assets at the same time, making use of their historical data before and after the crises and removing the general market noise to uncover their real behavioral traits under panic.

Besides, the research formulated a complex combination through the utilization of the Black-Litterman optimization tool in order to convert qualitative aspects to a mathematical environment. The new method managed to dispense with static risk inputs and instead use dynamic GARCH(1, 1) conditional volatility forecasts; also, it converted the past Shock Alpha scores to bounded, directional LLM sentiment signals on a simultaneous basis. This two-level architecture, representing the "Resilient" portfolio, gave the optimizer a chance to calibrate its mathematical confidence in accordance with the semantic context meaning that the model was not only able to inflict heavy penalties on fragile assets but also to allocate a large proportion of capital to structurally sound instruments during periods of high risk quite confidently.

The concluding part of the study served as a proof-of-work demonstrating that the theoretical model indeed worked in the real-world scenario. During the volatile 2024 test period, this architecture was empirically and econometrically validated out-of-sample. The empirical backtest, a simulation done through a tailor-made R-Shiny analytical application, unmistakably demonstrated the hybrid model's superiority. The Classical Markowitz portfolio showed extreme structural weakness. On the one hand, it suffered massive drawdowns; on the other hand, it severely limited capital growth due to its inflexible, backward-looking nature. The Agnostic control portfolio (BL + GARCH) was better from a risk-adjusted perspective and was able to show a better performance but it was not capable of fully unlocking the upside potential. On the contrary, the Resilient portfolio (BL + GARCH + LLM) showed the best absolute performance and the highest Sharpe ratio, which means it was able to build up wealth while at the same time dealing with market shocks.

Crucially, these performance metrics were corroborated by a battery of advanced econometric diagnostics. Factorial decomposition via OLS regression confirmed that the Resilient portfolio generated a statistically significant, positive Alpha, proving the strategy created genuine algorithmic value rather than merely assuming hidden systemic risk. Furthermore, Granger Causality testing proved that the Resilient portfolio was statistically immune to predictive drawdowns caused by broad market panics—a vulnerability that plagued the classical baseline. Finally, Bai-Perron structural break diagnostics confirmed that while the classical model suffered mathematical regime collapses during the acute crises of 2024, the dynamic Resilient portfolio maintained unbroken mathematical stability.

In summation, this Master's thesis successfully proves its core hypothesis. By mathematically equipping a Bayesian portfolio optimizer with the ability to dynamically forecast econometric risk and qualitatively differentiate between "fragile" and "anti-fragile" assets via LLM sentiment signals, it is possible to construct a portfolio that transcends traditional risk management. The resulting architecture provides institutional investors and capital allocators with a robust, empirically validated blueprint for surviving and thriving in an era of unprecedented informational volatility.

REFERENCES

1. Markowitz H. Portfolio Selection. *The Journal of Finance*. 1952. Vol. 7, No. 1. P. 77–91. DOI: 10.1111/j.1540-6261.1952.tb01525.x. URL: https://doi.org/10.1111/j.1540-6261.1952.tb01525.x.

2. Black F., Litterman R. Global Portfolio Optimization. *Financial Analysts Journal*. 1992. Vol. 48, No. 5. P. 28–43. DOI: 10.2469/faj.v48.n5.28. URL: https://www.tandfonline.com/doi/abs/10.2469/faj.v48.n5.28 .

3. Sharpe W. F. Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk. *The Journal of Finance*. 1964. Vol. 19, No. 3. P. 425–442. DOI: 10.1111/j.1540-6261.1964.tb02865.x. URL: https://doi.org/10.1111/j.1540-6261.1964.tb02865.x.

4. Fama E. F., French K. R. The Capital Asset Pricing Model: Theory and Evidence. *Journal of Economic Perspectives*. 2004. Vol. 18, No. 3. P. 25–46. DOI: 10.1257/0895330042162430. URL: https://doi.org/10.1257/0895330042162430.

5. Elton E. J., Gruber M. J. Estimating the Dependence Structure of Share Prices—Implications for Portfolio Selection. *The Journal of Finance*. 1973. Vol. 28, No. 5. P. 1203–1232. DOI: 10.1111/j.1540-6261.1973.tb01431.x. URL: https://doi.org/10.1111/j.1540-6261.1973.tb01431.x.

6. DeMiguel V., Garlappi L., Uppal R. Optimal Versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy? *The Review of Financial Studies*. 2009. Vol. 22, No. 5. P. 1915–1953. DOI: 10.1093/rfs/hhm075. URL: https://doi.org/10.1093/rfs/hhm075.

7. Kolm P. N., Tütüncü R., Fabozzi F. J. 60 Years of portfolio optimization: Practical challenges and current trends. *European Journal of Operational Research*. 2014. Vol. 234, No. 2. P. 356–371. DOI: 10.1016/j.ejor.2013.10.060. URL: https://doi.org/10.1016/j.ejor.2013.10.060.

8. Engle R. F. Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica*. 1982. Vol. 50, No. 4. P. 987–1007. DOI: 10.2307/1912773. URL: https://doi.org/10.2307/1912773.

9. Bollerslev T. Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics*. 1986. Vol. 31, No. 3. P. 307–327. DOI: 10.1016/0304-4076(86)90063-1. URL: https://doi.org/10.1016/0304-4076(86)90063-1.

10. Taleb N. N. Antifragile: Things That Gain from Disorder. New York : Random House, 2012. 519 p. ISBN: 9780812985924. URL: https://en.wikipedia.org/wiki/Antifragile_(book).

11. Bollerslev T., Chou R. Y., Kroner K. F. ARCH modeling in finance: A review of the theory and empirical evidence. *Journal of Econometrics*. 1992. Vol. 52, No. 1-2. P. 5–59. DOI: 10.1016/0304-4076(92)90064-X. URL: https://doi.org/10.1016/0304-4076(92)90064-X.

12. Engle R. F., Ng V. K. Measuring and Testing the Impact of News on Volatility. *The Journal of Finance*. 1993. Vol. 48, No. 5. P. 1749–1778. DOI: 10.1111/j.1540-6261.1993.tb05127.x. URL: https://doi.org/10.1111/j.1540-6261.1993.tb05127.x.

13. Ding Z., Granger C. W. J., Engle R. F. A long memory property of stock market returns and a new model. *Journal of Empirical Finance*. 1993. Vol. 1, No. 1. P. 83–106. DOI: 10.1016/0927-5398(93)90006-D. URL: https://doi.org/10.1016/0927-5398(93)90006-D.

14. Bollerslev T. The story of GARCH: A personal odyssey. *Journal of Econometrics*. 2023. Vol. 234, No. 1. P. 96–100. DOI: 10.1016/j.jeconom.2023.01.015. URL: https://doi.org/10.1016/j.jeconom.2023.01.015.

15. Andersen T. G., Bollerslev T., Diebold F. X., Labys P. Modeling and Forecasting Realized Volatility. *Econometrica*. 2003. Vol. 71, No. 2. P. 579–625. DOI: 10.1111/1468-0262.00418. URL: https://doi.org/10.1111/1468-0262.00418.

16. Lopez-Lira A., Tang Y. Can ChatGPT Forecast Stock Price Movements? Return Predictability and Large Language Models. *arXiv preprint arXiv:2304.07619*. 2023. URL: https://arxiv.org/abs/2304.07619.

17. Kirtac K., Germano G. Large language models in finance: what is financial sentiment? *arXiv preprint arXiv:2503.03612*. 2025. URL: https://arxiv.org/abs/2503.03612.

18. Lee Y., Kim Y., Kim J., Kim S., Lee Y. LLM-Enhanced Black-Litterman Portfolio Optimization. *arXiv preprint arXiv:2504.14345*. 2025. URL: https://arxiv.org/abs/2504.14345.

19. Mantshimuli L. A., Mwamba J. W. M. Enhancing portfolio optimization with multi-LLM sentiment aggregation: A Black-Litterman integration approach. *Investment Management and Financial Innovations*. 2025. Vol. 22, Issue 3. P. 213–226. DOI: 10.21511/imfi.22(3).2025.16. URL: https://doi.org/10.21511/imfi.22(3).2025.16.

20. Iacovides G., Konstantinidis T., Xu M., Mandic D. FinLlama: LLM-Based Financial Sentiment Analysis for Algorithmic Trading. *Proceedings of the 5th ACM International Conference on AI in Finance*. 2024. P. 134–141. DOI: 10.1145/3677052.3698656. URL: https://doi.org/10.1145/3677052.3698656.

21. Xing F. Designing Heterogeneous LLM Agents for Financial Sentiment Analysis. *arXiv preprint arXiv:2401.05799*. 2024. URL: https://arxiv.org/abs/2401.05799.

22. Kim S., Kim S., Kim Y., Park J., Kim S. LLMs Analyzing the Analysts: Do BERT and GPT Extract More Value from Financial Analyst Reports? *Proceedings of the Fourth ACM International Conference on AI in Finance*. 2023. P. 383–391. DOI: 10.1145/3604237.3626861. URL: https://doi.org/10.1145/3604237.3626861.

23. Ko H., Lee J. Can ChatGPT improve investment decisions? From a portfolio management perspective. *Finance Research Letters*. 2024. Vol. 64. P. 105433. DOI: 10.1016/j.frl.2024.105433. URL: https://doi.org/10.1016/j.frl.2024.105433.

24. Kumaran D., Fleming S. M., Markeeva L., Heyward J., Banino A. How Overconfidence in Initial Choices and Underconfidence Under Criticism Modulate Change of Mind in Large Language Models. *arXiv preprint arXiv:2507.03120*. 2025. URL: https://arxiv.org/abs/2507.03120.

25. Lo A. W., Ross J. Can ChatGPT plan your retirement?: Generative AI and financial advice. *Generative AI and Financial Advice*. 2024. URL: https://ssrn.com/abstract=4734895.

26. Вітлінський В. В., Великоіваненко Г. І. Ризикологія в економіці та підприємництві : монографія. Київ : КНЕУ, 2004. 480 с. ISBN: 966-574-569-7. URL: https://koha.oa.edu.ua/cgi-bin/koha/opac-search.pl?q=an:18769.

27. Камінський А. Б. Моделювання фінансових ризиків : монографія. Київ : Видавничо-поліграфічний центр «Київський університет», 2006. 303 с. ISBN: 966-594-840-7. URL: http://www.irbis-nbuv.gov.ua/cgi-bin/irbis_nbuv/cgiirbis_64.exe?Z21ID=&I21DBN=EC&P21DBN=EC&S21STN=1&S21REF=10&S21FMT=fullwebr&C21COM=S&S21CNR=20&S21P01=0&S21P02=0&S21P03=A=&S21COLORTERMS=1&S21STR=%D0%9A%D0%B0%D0%BC%D1%96%D0%BD%D1%81%D1%8C%D0%BA%D0%B8%D0%B9%20%D0%90$.

28. Примостка Л. О. Фінансовий менеджмент у банку : підручник. 3-тє вид., доп. і перероб. Київ : КНЕУ, 2012. 338 с. URL: https://buklib.net/books/21908/.

29. Барановський О. І. Фінансова безпека в Україні (методологія оцінки та механізми забезпечення) : монографія. Київ : КНТЕУ, 2004. 759 с. URL: https://lib.lntu.edu.ua/sites/default/files/2021-01/MONOGRAFIYA_dlya_zvitu_1.pdf.

30. Тяжкороб І., Сідельник О. Вплив геополітичних ризиків на інвестиційну активність банків з іноземним капіталом в Україні. *Фінансовий простір*. 2023. № 3 (51). С. 82–93. DOI: 10.30970/fp.3(51).2023.168177178. URL: https://doi.org/10.30970/fp.3(51).2023.168177178.

31. Чернєй В. В. Кримінально-правові та кримінологічні засади запобігання злочинам у сфері діяльності небанківських фінансових установ в Україні : монографія. Київ : КНЕУ, 2014. 456 с. URL: https://journals.maup.com.ua/index.php/economics/article/view/4664.

32. Кочетков В. М. Оцінка фінансової стійкості та стабільності банківської системи України на основі індикаторів. *Проблеми системного підходу в економіці*. 2019. № 6 (74), Ч. 3. С. 83–88. DOI: 10.32782/2520-2200/2019-6-59. URL: https://doi.org/10.32782/2520-2200/2019-6-59.

33. Galanos A. rugarch: Univariate GARCH models : R package version 1.5-5. 2026. URL: https://github.com/alexiosg/rugarch.

34. Ryan J. A., Ulrich J. M. quantmod: Quantitative Financial Modelling Framework : R package version 0.4.28. 2025. URL: https://github.com/joshuaulrich/quantmod.

35. Chang W., Cheng J., Allaire J., Sievert C., Schloerke B., Aden-Buie G., Xie Y., Allen J., McPherson J., Dipert A., Borges B. shiny: Web Application Framework for R : R package version 1.13.0. 2026. DOI: 10.32614/CRAN.package.shiny. URL: https://shiny.posit.co/.

36. Zeileis A., Leisch F., Hornik K., Kleiber C. strucchange: An R Package for Testing for Structural Change in Linear Regression Models. *Journal of Statistical Software*. 2002. Vol. 7, No. 2. P. 1–38. DOI: 10.18637/jss.v007.i02. URL: https://doi.org/10.18637/jss.v007.i02.

37. Pfaff B. VAR, SVAR and SVEC Models: Implementation Within R Package vars. *Journal of Statistical Software*. 2008. Vol. 27, No. 4. P. 1-32. DOI: 10.18637/jss.v027.i04. URL: https://doi.org/10.18637/jss.v027.i04.

38. R Core Team. R: A Language and Environment for Statistical Computing. Vienna, Austria : R Foundation for Statistical Computing, 2024. URL: https://www.R-project.org/.

39. Wolf T., Debut L., Sanh V., Chaumond J., Delangue C., Moi A., Cistac P., Rault T., Louf R., Funtowicz M., Brew J. HuggingFace's Transformers: State-of-

the-art Natural Language Processing. *arXiv preprint arXiv:1910.03771*. 2019. URL: https://arxiv.org/abs/1910.03771.

APPENDIX A

Table A.1. Matrix of Identified Informational Shocks (2021–2024) used for Shock Alpha Calibration

#	Date	Event Category	Event Description
1	2021-01-06	Geopolitical	US Capitol Attack
2	2021-01-28	Financial / Market	GameStop (GME) Short Squeeze
3	2021-02-15	Macroeconomic	Texas Power Crisis / Energy Shock
4	2021-03-11	Macroeconomic	Signing of the 'American Rescue Plan' ($1.9T)
5	2021-03-23	Macroeconomic	Suez Canal Blockage (Ever Given)
6	2021-05-12	Crypto-Specific	Tesla suspends Bitcoin vehicle purchases
7	2021-05-19	Crypto-Specific	China bans financial institutions from crypto services
8	2021-06-16	Monetary	Fed 'Tapering Talks' (FOMC Meeting)
9	2021-08-15	Geopolitical	Fall of Kabul (Afghanistan)
10	2021-09-20	Systemic Financial	Evergrande Liquidity Crisis Begins
11	2021-11-03	Monetary	Fed formally announces tapering of Quantitative Easing
12	2021-11-26	Macroeconomic	Identification of the COVID-19 'Omicron' variant
13	2022-02-24	Geopolitical	Russian Federation's full-scale invasion of Ukraine
14	2022-03-16	Monetary	First Federal Reserve Interest Rate Hike
15	2022-05-09	Crypto-Specific	Collapse of Terra (LUNA) algorithmic stablecoin ecosystem
16	2022-06-15	Monetary	Fed announces massive 75 basis point rate hike
17	2022-07-13	Macroeconomic	US CPI hits 9.1% (40-year inflation high)
18	2022-08-26	Monetary	Jerome Powell's Hawkish Jackson Hole Speech
19	2022-09-23	Systemic Financial	UK Mini-Budget Crisis (LDI pension fund shock)
20	2022-11-02	Monetary	Fed hikes 75 bps for the fourth consecutive time
21	2022-11-08	Crypto-Specific	FTX Cryptocurrency Exchange Bankruptcy
22	2023-01-10	Technological	Mainstream explosion/introduction of ChatGPT
23	2023-03-10	Systemic Financial	Silicon Valley Bank (SVB) Run and Failure
24	2023-03-12	Crypto-Specific	Signature Bank closure (Crypto banking crisis)
25	2023-03-19	Systemic Financial	Credit Suisse forced takeover by UBS
26	2023-06-05	Crypto-Specific	SEC files lawsuits against Binance and Coinbase
27	2023-07-26	Monetary	Fed raises rates to 5.25-5.50% (22-year high)
28	2023-08-17	Systemic Financial	Evergrande files for Chapter 15 US bankruptcy protection
29	2023-10-07	Geopolitical	Hamas Attack on Israel
30	2023-11-21	Crypto-Specific	Binance CEO CZ resigns and pleads guilty to US charges
31	2024-01-10	Crypto-Specific	SEC officially approves Spot Bitcoin ETFs

32	2024-02-22	Technological	Nvidia posts record financial results (AI Boom shock)
33	2024-04-13	Geopolitical	Iran launches unprecedented drone/missile attack on Israel
34	2024-04-19	Crypto-Specific	Fourth Bitcoin Halving Event
35	2024-07-11	Macroeconomic	US CPI drops unexpectedly, sparking massive sector rotation
36	2024-07-31	Monetary	Bank of Japan raises interest rates to 0.25%
37	2024-08-05	Systemic Financial	Global market sell-off & unwinding of the Yen Carry Trade

APPENDIX B

This appendix takes a detailed look into how the Large Language Model converts unstructured informational shocks into standardized quantitative signals. It starts with analyzing the "semantic gravity" of a news event and the further mapping of its structure changes onto a given asset.

The next table presents the conversion of two fundamentally different shock events: a systemic financial failure (negative shock) and a technological paradigm shift (positive shock).

Table B.1. Sample LLM Input/Output Logic for Semantic Signal Extraction

Feature	Case 1: Systemic Banking Crisis	Case 2: AI Technological Boom
Input Shock	SVB Collapse (March 2023): Sudden bank run and failure of Silicon Valley Bank due to interest rate exposure.	Nvidia GTC Keynote (March 2024): Announcement of the Blackwell GPU architecture and new AI industrial partnerships.
Target Asset	Bitcoin (BTC-USD)	Nvidia (NVDA)
LLM Assessment	"Event triggers extreme liquidity uncertainty and a flight from high-beta digital assets. High probability of contagion within the crypto-banking rails."	"Announcement confirms a sustained demand-side shock for hardware infrastructure. Demonstrates anti-fragile growth despite broad market volatility."
Shock Alpha Score	-0.85	+0.92

The scores generated show the model's skill at noticing the "fragile" type of volatility (where a drop in price actually indicates the company isn't doing well) from "resilient" volatility (where a price increase is actually a sign of good fundamentals). What makes it possible for the Bayesian optimizer to get the qualitative perspective (Q) that it needs from the portfolio for deciding to either hedge against a systemic failure or riding the technological momentum is that the LLM outputs these values in a bounded range.

APPENDIX C

In order to demonstrate the working of quantitative "volatility radar", the appendix shows the estimated GARCH(1, 1) coefficients for an important instrument in the portfolio in 2024 testing period. This model represents how risk stays over time and how quickly the market adjusts to fresh info.

Below is the table with the estimated parameters of Nvidia (NVDA) that correspond to its profile of a high-growth and high-volatility stock during the AI expansion cycle.

Table D.1. GARCH(1,1) Estimated Parameters for NVDA (Daily Data)

Parameter	Estimate	Standard Error	T-Statistic	Description
Ώ	0.000005	0.000001	4.82	Constant baseline variance
α	0.082	0.015	5.46	ARCH term: reaction to immediate shocks
β	0.895	0.018	49.72	GARCH term: volatility persistence (memory)

The high value of the β coefficient (0.895) indicates a significant "memory" effect, where volatility shocks persist over many trading days. The sum of α and β (0.977) is close to unity, which proves the presence of volatility clustering: large changes in returns tend to be followed by large changes, and small by small. For the Black-Litterman model, these coefficients mean that when a news event occurs, the GARCH engine produces a forecast that reflects not just today's panic, but a statistically grounded expectation of how long that risk will last, allowing for more precise capital allocation.