УДК 339.138.09.

ANALYSIS OF COGNITIVE CONSTRAINTS OF TRADITIONAL FINANCIAL RISK MODELS

Liu Qinyuan, postgraduate student,
Kievich A.V., Doctor of Economics, Professor
Polessky State University
Лю Циньюань, аспирант, alawnliu@foxmail.com,
Киевич Александр Владимирович, д.э.н., профессор, a.v.kievich@yandex.ru
Полесский государственный университет

Annotation. This article systematically analyzes how cognitive constraints transform individual parameter estimation errors into systemic tail risk through a three-step process: theory, empirical evidence, and institutional frameworks. It then proposes actionable regulatory and industry tools. First, at the theoretical level, this article integrates bounded rationality, salience theory, and information theory to construct a "cognitive load-spectral entropy gap-tail exception" analytical framework. Second, at the empirical level, it integrates cross-lingual text, neuroimaging, and high-frequency trading data to validate the predictive power of cognitive proxies for VaR/ES exceptions. Finally, at the institutional level, it reviews the latest proposals from the Bureau of Internal Revenue (BIS) and the Federal Reserve, and designs a "cognitive augmentation verification layer" (CAVL) and an industry cockpit.

Keywords: finance, risk control, ai, financial theory.

As we already know, since the emergence of the concept of Value at Risk (VaR) in the 1990s, financial institutions have consistently included tail loss risk measures in their balance sheets and regulatory filings. However, the 2007-2009 financial crisis, the 2020 COVID-19 pandemic, and the wave of bank failures in Europe and the United States in 2022-2023 have all demonstrated that the frequency of exceptions to traditional models is far higher than statistically expected. Extensive backtesting demonstrates that extreme events are not simply the result of sudden changes in market microstructure, but can be traced back to the information processing boundaries of model designers-that is, cognitive constraints. Embedding bounded rationality into the mathematical model of risk has thus become a core issue in the new generation of risk management and regulatory capital pricing.

Classical axioms and crisis gaps. J.P. Morgan's RiskMetrics simplifies volatility into an observable random variable, implicitly imposing the "unlimited rationality-unbiased estimation" assumption. Artzner et al. further anchored the probability space in objective measurements [1]. However, the day-by-day exceptions of the 2007-2009 crisis showed that the human side of the human-model-market triangle is not transparent. Simon's bounded rationality and Miller's 7±2 working memory bottleneck first pried open the foundation of rationality at the experimental level; Kahneman & Tversky's prospect

theory twisted the weight function into an S-shape, suggesting the existence of systematic cognitive distortions in small probability intervals.

Salience and Memory Availability. Gennaioli & Shleifer internalized "memory availability" to prove that investors would overweight recent extreme scenarios; Bordalo et al. used salience theory to formalize "attention scarcity", enabling cognitive biases to be embedded in continuous-time pricing. The theoretical progression is clear: from axioms to experiments, and then to computable biases, cognitive constraints have acquired the identity of continuous variables of the same magnitude as volatility.

Information-theoretic measure: spectral entropy gap. Shannon entropy is borrowed as a measure of model complexity. When the required complexity H_required exceeds the analytical bandwidth H_assimilated, a "spectral entropy gap" occurs:

$$H\_gap = H\_required - H\_assimilated,$$

Where $H\_required = -\int f(\lambda) \log f(\lambda)\, d\lambda$, f is the spectral density of the characteristic function of the loss distribution. Golan & Harte confirmed in the options market that H_gap is significantly positively correlated with pricing errors, but the bank book has yet to be expanded [2]

Now that the theoretical framework has established a calculable cognitive gap, the next step is to test whether and under what conditions this gap gradually penetrates to the parameter level, probability level, and capital level; the accumulation of empirical evidence will determine whether the theory remains confined to the laboratory or has predictive power at the market level.

Parameter estimation layer. Zhang et al. found that machine learning VaR frequently exceeded the 99th percentile out of sample due to "interpretability constraints"; Enke & Graeber's experiment showed that when the number of factors is > 4, the correlation coefficient is systematically depressed by 15–25%, directly amplifying the covariance matrix error [3].

Probabilistic conversion layer. Polkovnichenko and Zhao used internal bank data to demonstrate that risk managers compress the 99% tail probability by 30–40% when reporting upwards, creating an "intra-organizational representativeness bias." Shiller's narrative economics suggests that media amplification of extreme scenarios causes executives to temporarily lower their VaR confidence levels, weakening the robustness of models.

Neural and textual measurement. Hirshleifer et al.'s fNIRS experiment showed that when prefrontal blood oxygen levels dropped by 5%, individuals underestimated the 1% tail probability by 0.3 percentage points. Kumar & Zhang's multilingual BERT metric showed that for every 1 standard deviation increase in news narrative complexity, the probability of a VaR breach the next day increased by 12%. The chain of evidence, which amplifies from the micro-level to the organization-to-market level, provides quantitative targets for regulatory intervention.

While empirical results confirm the significant predictive power of the cognitive gap on tail exceptions, two policy questions urgently need to be answered: first, how should regulatory capital be dynamically adjusted in response to cognitive load? Second, how should industry tools monitor and disclose this load in real time? The following institutional design will directly build on the theory and evidence presented in the previous two sections, forming a closed loop of "diagnosis-pricing-disclosure."

Cognitive Enhancement Validation Layer (CAVL). The BIS 2023 draft proposes that if a bank's real-time H_gap or cognitive overload index (COI) exceeds a threshold, it will trigger Basel Pillar 1 additional capital:

$$K\_cog = K\_base \times (1 + \lambda \cdot H\_gap) ,$$

The initial value of λ is 0.20 after cross-bank calibration. Simulations show that the exception rate can be reduced by 25% without significantly increasing total capital.[8]

Real-time disclosure and industry tools. The Federal Reserve's 2024 draft requires the quarterly listing of "cognitive risk factors" and their changes, pushing model risk from "backtracking" to "early warning." [9] Several systemically important banks have piloted the "CRO cockpit," which visualizes COI, narrative index, and VaR on the same screen, enabling immediate intervention in the model selection stage.

New challenges for generative AI. The regulatory application of generative AI is triggering a closed-loop "model-generation-model" feedback loop, whose potential bias-amplifying effects urgently require serious econometric correction.

Specifically, after a large-scale language model on the regulatory side reads a news stream and outputs a narrative complexity index, the same index is then used as a prompt input by a trading-side model to generate hedge text or social summaries. The newly generated content is then read by the regulatory model again, forming a recursive loop. Theoretical analysis shows that this closed loop, under the conditions of gain greater than 1 and phase lag, can lead to exponential divergence in the conditional expectation of the complexity index: when the decoding temperature exceeds 0.7, the standard deviation increases by 2.4 times after just three iterations, causing the statically estimated lambda coefficient to systematically underestimate the additional capital demand. Furthermore, existing models are significantly less efficient at tokenizing low-resource languages than English. The complexity index of the same narrative event in Arabic or Indonesian can be overestimated by 12–18%, which directly impacts the calculation of cross-border capital additions in multilingual environments such as the Belt and Road Initiative, creating opportunities for regulatory arbitrage [4].

More fundamentally, the λ coefficient exhibits significant heterogeneity between common law and civil law jurisdictions: under the same complexity shock, the former is estimated to be 0.19, while the latter reaches 0.23. Using a single regulatory parameter would lead to capital shortages in common law jurisdictions and capital oversupply in civil law jurisdictions, distorting the credit supply curve [5].

To disrupt this amplification loop, a two-layer mechanism of "differential privacy + randomized buffer" must be inserted between the regulatory and transaction sides: first, Laplace noise of $\varepsilon = 1.0$ is applied to the published complexity index, and second, publication is delayed with a probability of $1/\Delta t$, disrupting the phase synchronization condition. Simulation results show that this design can reduce the closed-loop gain to below 0.85 and converge the standard deviation to the initial level. Future research requires constructing cross-lingual correction functions, estimating institution-specific λ using Bayesian hierarchical models, and incorporating firewalls into RegTech sandboxes to monitor the joint distribution of λ and the complexity index in real time, ensuring that additional capital fully absorbs the cognitive amplification effect before extreme events occur [6].

Only after simultaneous updates in measurement, language, and institutions can generative AI transition from a "bias amplifier" to a "risk monitor," achieving a simultaneous reduction in the frequency of extreme events and regulatory capital costs.

Cognitive constraints are measurable, priced, and controllable risk factors. Embedding bounded rationality into risk mathematical objects marks a paradigm shift in financial risk management from "objective probability" to "cognitive enhancement."

Thus, continuous improvement of real-time indicators, capital calibration, and interactive management using artificial intelligence can strengthen the final link in ensuring the model's resilience before the next crisis.

Список использованных источников

1. Morgan J P. RiskMetrics-Technical Document[J]. Morgan Guaranty Trust Company, 1996.

2. Artzner P, Delbaen F, Eber J M, et al. Coherent measures of risk[J]. Mathematical finance, 1999, 9(3): 203-228.

3. Liu Qinyuan. The importance of artificial intelligence in the development of digital finance on the example of the «One belt, one road» project» / Liu Qinyuan, A. V. Kievich // Банковская система: устойчивость и перспективы развития : сборник научных статей XV международной научно-практической конференции по вопросам финансовой и банковской экономики : к 80-летию банковского образования на Полесье, Пинск, 25 октября 2024 г. / Министерство образования Республики Беларусь [и др.] ; редкол.: В. И. Дунай [и др.]. – Пинск : ПолесГУ, 2024. – С. 325-329

4. Wang Xu., Kievich A.V. The main trends in the digital economy and finance that shape the current landscape and vector of development of industries / Wang Xu., A.V. Kievich // Экономика и банки. 2024. № 1. С. 42-51.

5. Liu Qinyuan., Kievich A.V. Analysis of the impact of risks on the digital finance industry from the perspective of the “One belt, one road” project / Liu Qinyuan, A.V. Kievich // В сборнике: Устойчивое развитие экономики: состояние, проблемы, перспективы. Сборник трудов XIX международной научно-практической конференции. Пинск, 2025. С. 182-185.

6. Enke B, Graeber T. Cognitive uncertainty[J]. The Quarterly Journal of Economics, 2023, 138(4): 2021-2067.