World Journal of Advanced Engineering Technology and Sciences

eISSN: 2582-8266

Cross Ref DOI: 10.30574/wjaets

Journal homepage: https://wjaets.com/

WJAETS

(Review Article)

Check for updates

Modernizing futures and options trading platforms for real-time margin and intraday risk computation

Swaminathan Vaidyanathan *

Sri Chandrasekharendra Saraswathi Viswa Mahavidyalaya, Enathur, Tamil Nadu, India.

World Journal of Advanced Engineering Technology and Sciences, 2026, 19(03), 054-063

Publication history: Received on 23 March 2026; revised on 26 May 2026; accepted on 28 May 2026

Article DOI: https://doi.org/10.30574/wjaets.2026.19.3.0254

Abstract

Modernizing futures and options trading platforms has become a strategic and technical priority because real-time margin and intraday risk computation must now operate under higher message rates, larger cross-product portfolios, and more demanding intraday control requirements. Traditional batch-based risk engines were designed for periodic recalculation, but current market conditions require exposures, sensitivities, and collateral requirements to be updated continuously in response to rapidly changing volatility and liquidity. Recent literature provides valuable building blocks for this transition, particularly in fast option pricing, surrogate modeling, deep-learning-based volatility estimation, reinforcement-learning-based hedging, and high-dimensional numerical methods. At the same time, the peer-reviewed evidence base remains fragmented, with far more attention to computational kernels than to full production-platform redesign for intraday margining. This literature review appraises recent journal articles in the context of platform modernization, such as computational latency, model fidelity, explainability and operational control. Major themes include event-based risk architecture, staged recomputation, rapid pricing of Greeks and exposure aggregation, and governance requirements for margin decisions that must remain defensible under stressed market conditions. The review identifies a persistent gap between advances in algorithms and deployable intraday margin systems, and argues that future progress will require integrated work across quantitative finance, distributed-systems engineering, and model-risk governance.

Keywords*:* Cloud-Native Risk Engines; Futures and Options Platforms; Intraday Risk; Real-Time Margin; Streaming Analytics

1. Introduction

Futures and options trading platforms are no longer being modernized merely for efficiency, but increasingly for control. In listed and cleared derivatives markets, conditions can change within seconds, yet many operational risk pipelines still reflect an earlier design in which end-of-day valuation, overnight exposure aggregation, and periodic collateral review were considered sufficient. When implied volatility shifts abruptly, correlation structures dislocate across contracts, or concentrated client positions generate nonlinear option risk, those changes must be translated into margin requirements almost immediately. End-to-end intraday margin platform redesign has not yet been covered extensively in recent journal literature; however, recent studies have already provided a number of enabling strands, particularly in the neural approximation of option value, faster pricing and hedging, and high-dimensional numerical techniques that reduce computation time without compromising accuracy [1].

The dynamic relationship between valuation and risk is the second driver of modernization. In futures and options trading, margin is not merely an accounting by-product. The logic of margin is based on an active model of current and stressed risk and cannot be decoupled from the computational techniques to update prices, Greeks, scenario losses and

* Corresponding author: Swaminathan Vaidyanathan.

netting effects in portfolios. The volatility models that are based on learning have fuelled this debate by demonstrating that the calibration and estimation of surface prices can be significantly expedited as compared to traditional repetitive optimization procedures [2]. For platforms that support intraday control, the central question is no longer whether fast approximation is possible, but whether it can be embedded in an auditable risk stack that ingests market events in real time and produces defensible collateral decisions.

Valuation of American-style and high-dimensional derivatives is of particular importance because the valuation of futures and options platforms are required to compute margin on portfolios with exercise features, nonlinear interactions among payoffs and cross-instrument hedging, which are too difficult to compute using brute-force simulation. Machine-learning-assisted pricing methods have shown that substantial computational gains are possible in high-dimensional settings, particularly through regression or surrogate approaches rather than repeated full revaluation [3]. This is directly relevant to intraday margin because any platform that cannot recalculate positions in near real time will struggle when market events require instruments to be revalued across multiple stress conditions. Platform modernization is not a cosmetic replacement of old software with new software, but a re-architecture of the valuation-to-collateral pipeline around selective recomputation, state retention, and model choice under latency constraints.

The other theme of recent literature is connected with the operational use of machine learning beyond pricing. Reinforcement-learning-based hedging studies suggest that effective trading and risk-response policies can be learned dynamically in the presence of transaction costs and market frictions [4]. Despite the separation between hedging and margin, they are both based on the need to update states rapidly, represent risk in a path-sensitive fashion, and have the ability to react to intraday market changes without destabilizing the system by the introduction of excessive noise or model opaqueness. At the same time, continuous numerical simulations of high-dimensional nonlinear partial differential equations have expanded the range of computational techniques that can be employed to implement a risk engine in real time, especially in those applications where finite-difference methods or lattice methods become slow or too delicate to be employed to model a risk engine in real time [5]. All of these developments are contributing to a gradual re-definition of the potential design space of intraday derivatives risk infrastructure.

Despite these advances, there are major issues to be addressed. Peer-reviewed literature has been much more interested in individual computational kernels than the design of whole futures and options trading systems that can receive live market data, recalculate sensitivities, and aggregate portfolio risk and issue or update margin requirements subject to governance constraints. Issues related to the explainability, rollback control, stream consistency, model fallback, and operational resilience are still underdeveloped in the recent journal literature. This review addresses that gap by synthesizing evidence published between 2020 and 2025 on modernization of futures and options trading platforms for real-time margin and intraday risk computation. The sections below comment on the recent literature, systematize the main methodological tools, compare the reported results through a platform-engineering lens and describe the long-term limitations, and research agenda of the next stage of modernization of intraday risk.

2. Literature Review

The recent journal literature relevant to real-time margin and intraday risk modernization can be organized into four main groups. The first cluster concerns fast surrogate valuation and hedging, in which neural approximators replace repeatedly expensive pricing procedures with trained mappings from state variables to prices or sensitivities. The second category is the high-dimensional numerical methods, especially deep schemes of partial differential equations and optimal stopping, which are of significance in option-heavy portfolios, that are difficult to compute using conventional methods. The third cluster is related to volatility and market-model representation, with deep calibration and neural market models which seek to retain arbitrage structure and speed up model updates. The fourth group concerns decision-oriented control, especially reinforcement-learning-based hedging, and helps explain how stateful intraday risk decisions might be computed in real time. The literature of these clusters has a wealth of computational innovation and a paucity of end-to-end platform research.

The surrogate-pricing literature is especially important because real-time margining is constrained first by computation rather than by economic logic alone. Ruf and Wang defined pricing and hedging based on neural networks as a general emerging literature, but also highlighted the reliance on the design of training, data visualization and extrapolation beyond the training scope [1]. Goudenège, Molent, and Zanette showed that high-dimensional American-option pricing can be solved with machine learning in the regimes where traditional methods are either costly or unstable [3]. Lapeyre and Lelong have extended this direction to Bermudan option pricing which shows that learning structures based on regression can be useful in the exercise-dependent valuation [7]. One implication of this work is that there can be no single model of intraday platform modernization, because latency tolerance, instrument complexity, and auditability requirements differ across risk tasks.

A second body of literature is about numerical schemes which could be scaled to the state dimensionality of modern options books. Huré, Pham and Warin constructed deep backward schemes of high-dimensional nonlinear partial differential equations and Beck and others came up with deep splitting of the parabolic PDEs [5], [8]. Becker and co-authors proceeded to create high-dimensional optimal stopping methods based on deep-learning [6]. These are not platform papers in the strict engineering sense, but they are directly relevant to the present topic: futures and options platforms have to be modernized to use computational primitives that can quickly recalculate when there is a scenario shock. Real-time margin is an architectural dream instead of a working control mechanism, without scalable numerical kernels.

The volatility and market-model literature present another, yet equally important factor. Deep-learning-based volatility modelling can greatly improve the speed and flexibility of the calibration, which is significant in options books as the margin and intraday risk of the book are dependent on the current shape of the surface and not a volatility point estimate as shown by Horvath, Muguruza, and Tomas [2]. Bayer and co-authors went to even greater lengths and calibrated rough stochastic-volatility models using deep learning [9]. The control of arbitrages was brought closer to the core ideas of Cohen, Reisinger and Wang in neural-SDE market models which were intended to be fully consistent with the financial structure and which had the advantage of the flexibility of learning [10]. This literature matters because margin engines are highly sensitive to model inconsistency: a fast intraday engine is of limited value if its surface representation produces unstable or economically unsound loss scenarios.

The decision-oriented literature is smaller but has conceptual importance. Applying reinforcement learning to dynamic hedging, Cao and co-authors discovered that sequential decision models can capture the internalization of transaction costs and path dependence, which cannot be approximated using the simplified hedge [4]. From the perspective of platform modernization, this implies that the intraday risk control can be changed to continuous feedback instead of periodic measurements. Even so, a sharp gap remains between the literature and the production requirements of listed-derivatives platforms. Recent articles rarely discuss stateful stream processing, Greek cache-invalidation policies, low-latency Greeks aggregation of margin factors on large books of clients or model governance of live intraday margin calls. The literature thus has good contents to be used in the modernization, but few studies on the integrated platform.

Table 1 summarizes the recent literature most directly relevant to the present topic. The trend is prominent. Nearly all the studies add some significant computational or modelling aspect, but few take the entire chain of market-data event to intraday margin decision. This fragmentation is itself an important finding, since it shows that the academic discipline is yet to converge on the platform level issue that exchanges, brokers, and clearing-facing infrastructures need to address.

Table 1 Summary of key findings

Ref	Focus	Key Findings
[1]	Neural-network approximation for option pricing and hedging	Identified fast surrogate valuation and hedge mapping as a credible alternative to repeated full revaluation, while stressing dependence on training-domain coverage and out-of-sample robustness.
[2]	Deep-learning volatility modelling for derivatives	Showed that learned volatility representations can improve calibration speed and flexibility, with implications for intraday repricing of option books and margin shock sensitivity.
[3]	Machine learning for American-option pricing in high dimensions	Demonstrated that learning-based approximators can materially reduce computational burden in exercise-sensitive portfolios relevant to intraday margin recalculation.
[4]	Reinforcement-learning-based dynamic hedging	Reported that sequential learning policies can adapt hedges under costs and frictions, highlighting continuous-control logic relevant to intraday risk governance.
[5]	Deep backward schemes for high-dimensional nonlinear PDEs	Established a scalable route for solving nonlinear pricing and risk equations that would otherwise strain real-time platform budgets.

[6]	Deep learning for high-dimensional optimal stopping	Showed that optimal stopping problems can be handled in dimensions relevant to complex derivatives books, supporting rapid recomputation for exercise-sensitive risk.
[7]	Neural-network regression for Bermudan option pricing	Provided evidence that regression-based approximators can support repeated valuation of callable or exercise-linked products with lower latency.
[8]	Deep splitting for parabolic PDEs	Contributed a decomposition strategy that improves tractability of complex valuation problems, making selective intraday recomputation more realistic.
[9]	Deep calibration of rough stochastic-volatility models	Indicated that richer volatility dynamics can be calibrated faster, reducing the tension between model fidelity and operational timeliness in option platforms.
[10]	Arbitrage-free neural-SDE market models	Addressed the need for structural coherence in learned market models, a prerequisite for defensible intraday margin and scenario-based loss estimation.

3. Methodology

The literature reviewed here is best synthesized through a layered control framework rather than a single-model framework. The calculations of real-time margin and intraday risk on futures and options exchanges today are dependent on the interaction of five layers: market-event ingestion, pricing and sensitivity computation, stateful portfolio aggregation, margin-policy translation and governance or fallback control. In more recent journal papers, not all five layers are usually taken into consideration, but only one or two. Computational kernels are provided by pricing and hedging papers; a fast state estimation is provided by volatility papers; scalability is provided by numerical-method papers; adaptive control logic is provided by decision-oriented work. Only a coherent operational design can combine these layers and make platform modernization possible.

Figure 1 presents a topic-specific framework for understanding this interaction. The figure is not intended as a generic input-process-output flowchart. It is organized around the actual decision path implied by the topic: market events update a state store, the state store triggers selective repricing and Greeks refresh, those analytics feed exposure aggregation, and the aggregated outputs feed margin and alert logic; a governance layer checks or overrides live outputs in the event of instability, latency breaches, or model drift. This model reflects one of the significant aspects of the literature: the binding constraint on intraday margin is frequently not the accuracy of valuation per se, but the coordination of accuracy, latency and control in the face of ever-evolving state.

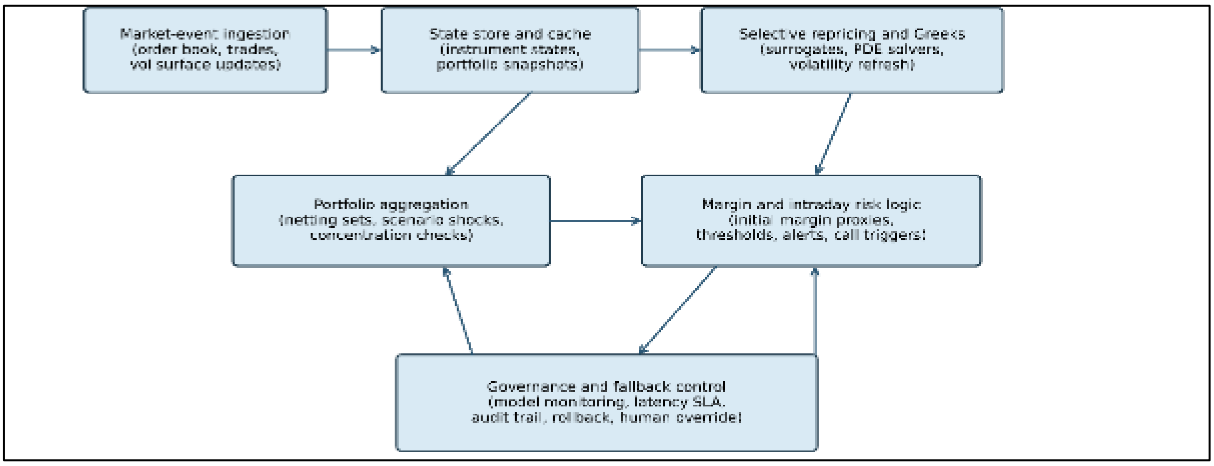

Figure 1 Layered framework for real-time margin and intraday risk computation on modernized futures and options platforms

The literature can be classified in four operational families in terms of methodology. The first family may be described as direct surrogate modeling, in which neural networks or regressors are used to predict prices, Greeks, continuation values, or scenario losses [1], [3], [7]. The second family comprises deep numerical solutions for PDE and optimal-stopping problems, especially where instrument complexity or dimensionality makes conventional methods too slow for practical use [5], [6], [8]. The third family is fast calibration and market-state representation, like learned volatility surfaces and market models that are mindful of arbitrage [2], [9], [10]. The fourth family is adaptive control which was represented by the reinforcement-learning-based hedging which shows how intraday risk behavior might be associated with how the market is evolving progressively rather than the snapshots [4].

Each family has distinct strengths and limitations when evaluated from a platform-modernization perspective. The raw reduction in latency is maximum in the case of surrogate models, and they are more prone to regime shift and less interpretable. Deep numerical methods are more in touch with structural finance models, but are still often computationally expensive unless used with care as selective recomputation kernels as opposed to universal engines. The learned volatility and neural-SDE models are more reflective of the state, however, there must be a platform which can version, monitor and roll back state representations when anomalous behavior occurs. Ongoing decision support during the day is prospective with adaptive control methods, yet live margin systems demand greater auditability than is often provided with reinforcement-learning studies. A credible modernization agenda therefore requires method portfolios, layered validation, and contingency paths rather than exclusive reliance on a single computational paradigm.

4. Discussion

The reviewed studies support a strong but cautious conclusion: meaningful advances in real-time margin and intraday risk computation are technically achievable, but current journal evidence is much stronger on component acceleration than on integrated production deployment. Strong evidence lies in the computational heart of the problem. Neural surrogates, deep regression architectures, and high-dimensional numerical algorithms have repeatedly demonstrated that option pricing and other valuation problems can be sped up compared to standard brute-force algorithms [1], [3], [5]. This has relevance in the modernization of futures and options platforms in the sense that the value of repeated valuation under shocks, concentration changes, and cross-instrument aggregation is ultimately constrained by margining speed. The literature therefore helps to alleviate a major barrier, but by itself, does not address platform-level control.

A second recurring pattern concerns the latency-fidelity trade-off. Surrogate models and fast regression models are typically the lowest-latency models after training, and therefore are deployed in live pipelines in intraday. These advantages, however, presuppose a certain stability in the field, training coverage, and careful treatment of extrapolation. Structural contact can be extended to the financial theory in deep numerical schemes as well as optimal-stopping schemes, which can be more realistic in case of regime change, but have more expensive computational profile, and are therefore more difficult to implement in high-frequency intraday refresh loops [5]-[8]. The implication of the platform is self-evident: modernization cannot be achieved by selecting a single best method in isolation. It includes assigning different techniques to different degrees of service such as sub-second notifications, minute re-calculations, slower validation runs.

Figure 2 depicts this trade-off by positioning method families based on the literature. The figure demonstrates that direct surrogates and learned volatility updates are located in the low-latency area with the deep PDE and optimal-stopping algorithms being located at higher fidelity but increased computational cost. This supports the argument that real-time margin systems should be stratified, with very fast approximate engines linked to slower but structurally robust validation layers.

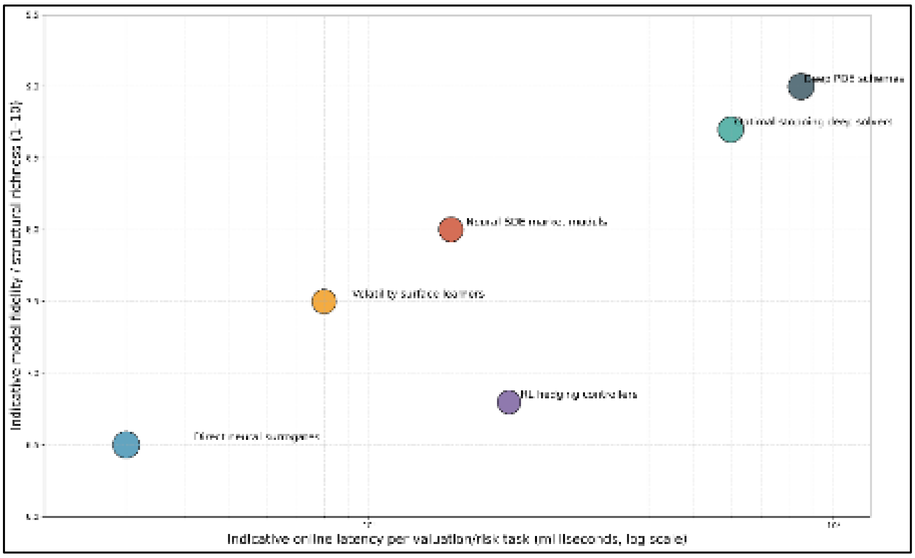

Figure 2 Literature-derived latency–fidelity trade-off across computational approaches

The third trend is in relation to the state representation, particularly of options books. Horvath, Muguruza, and Tomas show that deep learning can improve the speed and flexibility of volatility-model calibration, and calibrate more quickly [2], and Bayer et al. extended the idea of fast calibration to rough-volatility conditions [9]. The structural aspect of the same problem was dealt with by Cohen, Reisinger and Wang who focused on arbitrage-conscious neural market models [10]. These contributions are more consequential than can be implied by a narrow reading, in the case of intraday margin computation. The encoding of the state of the market is very sensitive to margin. A platform that updates prices quickly using an unstable or economically inconsistent surface representation may produce intraday margin movements that are fast but unreliable. It is therefore shown in the literature that compute speed must be matched by state quality.

The reviewed evidence also shows that exercise-sensitive instruments remain one of the clearest stress points in intraday risk computation. The applicability of American and Bermudan option techniques goes way beyond the niche of such a type of product because the reasoning of exercise is a proxy for the more universal problem of path dependence and nonlinear decision frontiers in a portfolio [3], [6], [7]. A re-architectured real-time risk platform should thus have the capability of dealing with heterogeneous computational kernels as opposed to putting all products through one repricing mechanism. This requirement pushes platform architecture toward service decomposition, cache-aware portfolio segmentation, and selective invalidation based on staleness sensitivity. There is a great deal of literature on such engineering decisions in peer-reviewed journals, but the literature on computational issues strongly indicates that such decisions are necessary.

Table 2 contrasts the main methodological strategies of the literature reviewed from the perspective of platform modernization. The comparison reveals that there is no method which is consistently dominant. Neural surrogates are strong operationally in terms of recalculation many times, though the literature often warns of sensitivity to domain shifts. Deep PDE and optimal-stopping methods can be used to improve structural defensibility, but must be carefully coordinated not to ruin latency targets. Volatility and neural-SDE methods help increase the coherence of states, but necessitate version control and tracking of arbitrage. Reinforcement learning offers adaptive control logic and current studies are yet to tackle the documentation and override requirements that are expected in live margin system.

Table 2 Method comparison

Ref	Method	Strengths	Limitations
[1]	Neural-network surrogate pricing and hedging	Very low online latency after training; supports repeated Greeks and scenario updates	Performance tied to training domain; difficult error attribution during regime breaks
[2]	Deep-learning volatility modelling	Fast surface refresh; flexible nonlinear representation of smile and term structure	Surface stability and economic interpretability require additional control layers

[3]	Machine-learning approximation for American-option valuation	Handles high-dimensional exercise-sensitive instruments more tractably than brute-force methods	Continuation-value approximation error can propagate into scenario loss estimates
[4]	Reinforcement-learning hedging	Captures sequential control under costs and frictions; useful for intraday adaptive logic	Auditability and policy explainability remain weak for live collateral decision contexts
[5]	Deep backward schemes for nonlinear PDEs	Scales to dimensions that obstruct classical methods; relevant to complex risk equations	Training and convergence management can be computationally demanding in production settings
[6]	Deep optimal stopping	Strong fit for callable and exercise-dependent structures in option-heavy portfolios	Online deployment still heavier than simple surrogates for ultra-fast alert loops
[7]	Neural-network regression for Bermudan pricing	Supports repeated evaluation of exercise-linked products with reduced computation time	Product-specific training can limit broad portfolio portability
[8]	Deep splitting for parabolic PDEs	Decomposes difficult valuation problems into computationally manageable components	Integration into live streaming risk stacks remains underexplored in journal literature
[9]	Deep calibration of rough-volatility models	Preserves richer dynamics while accelerating calibration cycles	Model richness raises governance burden and parameter-monitoring complexity
[10]	Arbitrage-free neural-SDE market models	Improves structural coherence of learned state representation for risk and margin use	Implementation complexity and calibration workflow may strain operational simplicity

This comparison leads to the operational suitability map shown in Figure 3. The narrative is criterion-based (not bibliometric) in nature because the key question is not a number of publications but platform fit. The figure suggests that direct surrogates tend to perform better on low latency and ease of integration, arbitrage-conscious market models and deep PDE methods will tend to be more successful in high structural consistency and stress defensibility, and the case that multi-tier risk stacks, rather than single-engine systems, are more likely to succeed.

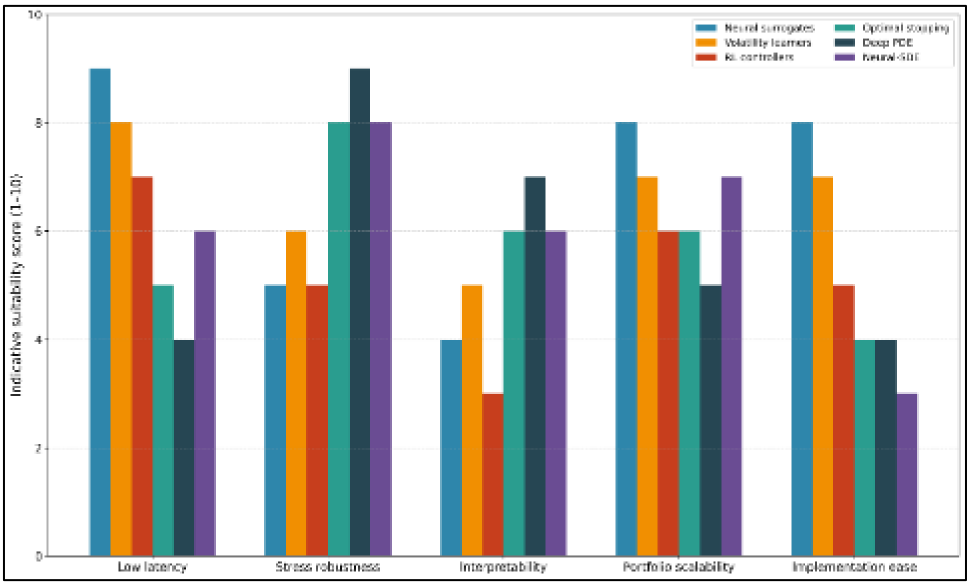

Figure 3 Operational suitability of method families for intraday margin platforms

Table 3 presents the literature reported to the system level through comparison of the nature of the outcomes that are valued to the title topic: speed, dimensional scalability, calibration quality, structural coherence and policy adaptiveness. One of the limitations shared by the studies reviewed can be observed with the help of the table. Most of the reported findings have to do with the performance of the algorithmic or numerical performance alone. Metrics such as stale-risk frequency, intraday margin false positives, rollback rates, cache-invalidation efficiency, and governance-intervention frequency are reported in very few studies to provide platform-relevant measures. The lack of it is among the research gaps which have been identified as the strongest in this review.

Table 3 Results comparison

Ref	System / Context	Metric / Basis of Comparison	Outcome
[1]	Option pricing and hedging approximators	Online valuation and hedge computation speed versus classical revaluation	Reported major acceleration potential, with accuracy dependent on training coverage and target function complexity
[2]	Learned volatility modelling for derivatives	Calibration speed and surface flexibility	Reported faster model adaptation with richer nonlinear surface representation
[3]	High-dimensional American-option setting	Pricing tractability and approximation quality in large state spaces	Showed feasible evaluation of exercise-sensitive products beyond the comfort zone of standard methods
[4]	Dynamic hedging under market frictions	Sequential policy performance under transaction costs	Reported adaptive hedge behavior superior to static-rule baselines in the studied settings
[5]	High-dimensional nonlinear PDEs	Scalability of numerical solution procedure	Demonstrated viability of deep schemes where conventional solvers become computationally burdensome
[6]	Optimal stopping problems relevant to early exercise	Dimensional scalability and value approximation	Reported strong performance in problems that are central to option-heavy risk books
[7]	Bermudan option pricing	Regression efficiency and repeated valuation capability	Indicated that trained regressors can support recurring valuation with reduced computational overhead
[8]	Parabolic PDE valuation tasks	Decomposition-based computational tractability	Showed that splitting can make complex pricing tasks more manageable for repeated use
[9]	Rough-volatility calibration workflows	Speed of calibration with richer volatility dynamics	Suggested that structural richness need not be sacrificed entirely for faster calibration cycles
[10]	Learned market-state models with arbitrage control	Economic coherence of generated market dynamics	Addressed a key prerequisite for trustworthy intraday scenario and margin computation

A final synthesis organizes the reviewed literature through an application-oriented control timeline. It demonstrates how a modernized platform is able to simulate different types of risk calculations in very short, short and supervisory intraday horizons. This value plays a role in one of the key theses of the paper: the modernization of the real-time margin cannot be perceived as the constant recomputation of all values after each event. Tiers are a more reasonable design with micro-latency approximate updates made, periodic validation checks, and governance checkpoints to protect against silent model drift or unstable margin output.

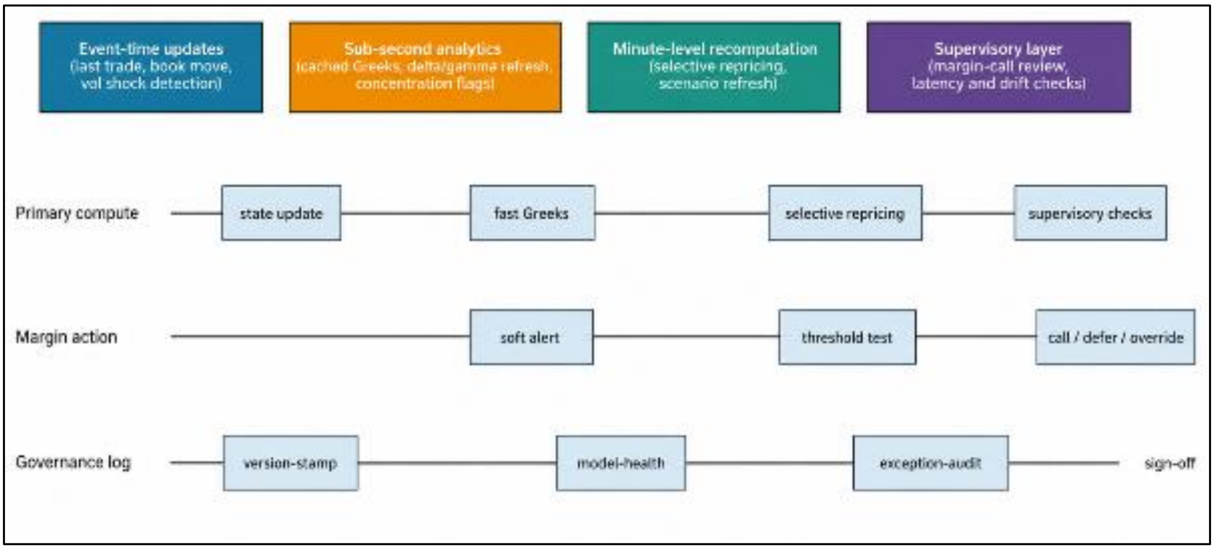

Figure 4 Tiered intraday control timeline for real-time margin computation

Collectively, the studies reported are in support of a platform architecture, which is tiering, selectivity, and governed approximation. The literature is not in support of an ambitious but naive attempt to re-compute full-book full-scenario full-fidelity risk at every tick. It supports a more disciplined modernization program in which most event flow is handled by approximate analytics, selective recomputation is triggered by material state changes and the governance layer constantly measures whether the fast path remains within pre-specified tolerances. The inherent paradox of the literature is that computational ambitions and operational realism are in tension: the research papers are biased to achieve as much numerical strength as possible, and live intraday margin systems are biased to achieve as much speed as possible, stability as well as explainability and recoverability. This contradiction appears to define the current stage of the field.

5. Future Directions

The most urgent research priority is the development of journal-quality studies focused on complete intraday risk pipelines rather than isolated pricing components. The recent literature has shown conclusively that computational friction can be alleviated through fast valuation, learned volatility representation as well as high-dimensional numerical methods. What is yet to be developed is the question at the system level of how these elements need to be coordinated within live futures and options platforms. Future research ought to hence document architecture-conscious measures like event-to-margin latency, stale-state rate, fallback activation rate, amplification of margins under market pressure and explainability of operation. In the absence of such measures, the field will continue to produce algorithmic advances while underreporting what actually matters for deployable intraday control.

The second priority has to do with differentiable and decomposable margin logic. The current body of literature is predisposed to either heading to pricing, hedging or calibration, when production margin systems have to combine sensitivities, stress cases, concentration add-ons, offsets or governance thresholds. Model designs that distinguish fast differentiable components from slower rule-based overlays would be valuable. These decompositions may enable high frequency refresh of the economically smooth component of the risk and hand over more expensive supervisory logic to be called at checkpoint intervals or events. This would take the computational design much nearer to the actual form of operation of intraday margining on modernized listed-derivatives platforms.

A third priority concerns model governance for fast-learning systems. Real-time margin is a high-consequence output. A fast engine that cannot explain an atypical intraday margin jump will struggle to remain viable in production. Future journal work should therefore integrate quantitative finance, model risk management, and distributed-systems reliability. Uncertainty tagging of surrogate outputs, online drift diagnostics of volatility learners, learned model layered fallback to conservative benchmark calculators and post-event audit of versioned replay environments are other important areas. Progress in this regard would help to narrow the gap between the impressive numeric literature and the vastly superior levels of live infrastructure of the markets.

A fourth priority is interdisciplinary benchmarking. The field needs standardized benchmark exercises that are anchored on modernization of the futures and options platforms rather than the futures and options platform independent valuation exercises. Examples of benchmarks of interest would include large option portfolios with cross-expiry dependencies, intraday volatility shocks, concentrated positions and service level constraints, as well as realistic constraints. The competing methods performance could subsequently be compared not only with regard to price error, but also with regard to latency, resilience, explainability, and margin stability. Such research, one would be able to compare direct surrogates and deep PDE methods, neural-SDE market models and adaptive controllers on the grounds that most critical to the title topic, i.e., the ability of a platform to compute intraday risk and margin in real time without compromising control.

6. Conclusion

Modernizing futures and options trading platforms for real-time margin and intraday risk computation is fundamentally a problem of controlled computation under live-market uncertainty. The recent journal literature has provided solid evidence that the key computational bottlenecks can be mitigated with the help of neural surrogates, high-dimensional deep numerical schemes, fast volatility learning, and adaptive control. These developments materially expand what is feasible in intraday derivatives-risk infrastructure.

However, the literature remains uneven. Most recent research considers valuation, calibration, or hedging aspects independently, whereas production-level intraday margin involves the coordination of market-event ingestion, stateful portfolio aggregation, selective recomputation, margin-policy logic and governance. The most important unresolved issue is the gap between algorithmic performance and platform-level deployment. Fast analytics are not sufficient to form a modernized platform without managing latency, explainability, rollback and margin stability.

The major academic implication is that future developments will be based on the incorporation of computational finance into systems engineering and model governance. The field has already demonstrated that real-time calculation is increasingly possible. The next task is to ensure that computation is reliable, interpretable, and operationally defensible in live futures and options platforms, where intraday volatility, nonlinear risk, and collateral decisions cannot wait for end-of-day batch processing.

References

[1] Ruf, J., & Wang, W. (2020). Neural networks for option pricing and hedging: a literature review. Journal of Computational Finance, 24(1), 1–46.

[2] Horvath, B., Muguruza, A., & Tomas, M. (2021). Deep learning volatility. Quantitative Finance, 21(1), 11–27.

[3] Goudenège, L., Molent, A., & Zanette, A. (2020). Machine learning for pricing American options in high-dimensional Markovian and non-Markovian models. Quantitative Finance, 20(4), 573–591.

[4] Cao, J., Chen, J., Hull, J., & Poulos, Z. (2021). Deep hedging of derivatives using reinforcement learning. Journal of Financial Data Science, 3(1), 10–27.

[5] Huré, C., Pham, H., & Warin, X. (2020). Deep backward schemes for high-dimensional nonlinear PDEs. Mathematics of Computation, 89(324), 1547–1579.

[6] Becker, S., Cheridito, P., Jentzen, A., & Welti, T. (2021). Solving high-dimensional optimal stopping problems using deep learning. European Journal of Applied Mathematics, 32(3), 470–514.

[7] Lapeyre, B., & Lelong, J. (2021). Neural network regression for Bermudan option pricing. Monte Carlo Methods and Applications, 27(3), 227–247.

[8] Beck, C., Becker, S., Cheridito, P., Jentzen, A., & Neufeld, A. (2021). Deep splitting method for parabolic PDEs. SIAM Journal on Scientific Computing, 43(5), A3135–A3157.

[9] Bayer, C., Horvath, B., Muguruza, A., Stemper, B., & Tomas, M. (2021). On deep calibration of rough stochastic volatility models. Quantitative Finance, 21(11), 1831–1845.

[10] Cohen, S. N., Reisinger, C., & Wang, S. (2021). Arbitrage-free neural-SDE market models. Applied Mathematical Finance, 28(5), 417–450.