

Building an adaptive tokenomics model using reinforcement learning algorithms to ensure long-term economic sustainability of crypto projects.

Sergei Vertepov

Senior Backend Engineer, Faraway LLC
Yerevan, Armenia

Abstract

The article examines the robustness of algorithmic tokenomics models under conditions of extreme volatility in cryptocurrency markets, using the collapse of the Terra/Luna ecosystem (2022) and liquidity crises in 2019–2023 as empirical material. It is shown that static, deterministic monetary supply control loops exhibit a systemic vulnerability: under abrupt shifts in market conditions, fixed rules for issuance and liquidity withdrawal lag, do not account for the distribution of scenarios, and incorrectly process nonlinearities in stress regimes, which increases the probability of the system entering unstable trajectories. As an alternative, an adaptive economic architecture is substantiated, in which regulatory functions are implemented by an intelligent agent based on deep reinforcement learning (DRL). Proximal Policy Optimization (PPO) is selected as the baseline algorithm, ensuring stable policy updates under an optimization step constraint, which is critical for preventing cascading destabilization. The state space and action set of the regulator agent are formalized, enabling the dynamics of key macro-parameters of tokenomics to be described as a controlled stochastic process. Simulations in cadCAD demonstrate the advantage of DRL stochastic policies compared to PID controllers in maintaining the peg and supporting liquidity under conditions of partial observability and noisy perturbations. Additionally, current requirements of legal regulation of digital financial assets are taken into account, and approaches are proposed for integrating DRL agents into smart contracts with on-chain execution, controlled access to data, and preservation of verifiability and reproducibility of the mechanism.

Keywords: adaptive tokenomics, algorithmic stablecoins, stochastic control, deep reinforcement learning, Proximal Policy Optimization, cadCAD simulations, on-chain execution of smart contracts.

Introduction

The rapid development of decentralized finance (DeFi) and the market of digital financial assets (DFAs) in 2019–2023 has made the task of maintaining macroeconomic stability within closed tokenized loops relevant for economic theory and for the engineering of financial protocols. Tokenomics schemes based on rigidly hard-coded algorithmic rules for token issuance and burning have demonstrated limited viability under exogenous shocks and targeted speculative impacts, because they relied on the assumption of a relatively stationary market environment and the predictability of participant responses [1, 11].

This deficit of robustness manifested most clearly in the class of algorithmic stablecoins. The collapse of the Terra ecosystem (UST/LUNA) in May 2022, which led to a loss of market capitalization exceeding USD 40 billion, has become established in the literature as a characteristic case of a death spiral: the loss of confidence by economic agents initiated a self-reinforcing liquidity outflow, while the static arbitrage algorithm proved unable to restructure its operating regime amid an avalanche-like growth of imbalances [3]. Empirical and modeling studies indicate that mechanisms showing apparent effectiveness in growth phases can, when the market cycle changes, shift into a mode of accelerated system destruction,

because they do not incorporate behavioral components of decision-making and do not account for nonlinear relationships between market depth and volatility [5, 15].

In parallel, progress in artificial intelligence, primarily in reinforcement learning (RL) methods, has formed a technological basis for constructing self-regulating economic systems. In particular, the AI Economist concept demonstrates the applicability of two-level RL to optimizing parameters of tax policy and redistribution mechanisms, showing results that outperform traditional theoretical constructs of classical economics in terms of efficiency and adaptivity [7, 24]. Transferring a similar paradigm to tokenomics makes it possible to shift the focus from reactive adjustments to predictive control, in which an intelligent agent continuously refines a strategy for maximizing the target function of system stability under conditions of stochastic uncertainty and partial market observability [1, 13, 16].

Within the scientific discussion, it is emphasized that smart contracts should be considered not only as program code, but also as a legal-technical mechanism for the performance of obligations, which increases the requirements for the predictability and verifiability of the regulatory rules embedded in them [9]. Under contemporary conditions of limited availability of traditional funding channels, DFAs are often positioned as an alternative instrument for raising capital and investing, which strengthens the demand for reliability and economic security of the issued assets [10, 17]. In this logic, the introduction of adaptive algorithms capable of automatically stabilizing value and liquidity is considered a potentially significant factor in strengthening trust in Russian DFA platforms.

The central **objective** of the study is to construct a theoretical and methodological foundation for an adaptive tokenomics model based on the Proximal Policy Optimization (PPO) algorithm, oriented toward maintaining exchange-rate stability and managed asset liquidity under stochastic disturbances. To achieve this objective, the study provides for a retrospective analysis of the causes of destabilizations of algorithmic stablecoins with identification of the limitations of deterministic designs; formalization of the tokenomics control problem as a Markov decision process (MDP), enabling a rigorous specification of the spaces of states, actions, and rewards; justification of the choice of PPO as the method for training the regulator agent with regard to requirements for training stability and control of policy updates; simulation stress testing in the cadCAD environment on historical data from crisis periods; and a comparative assessment of the effectiveness of the proposed model relative to baseline algorithmic strategies.

The scientific novelty consists in substantiating the transition from static deterministic issuance and liquidity withdrawal loops to a stochastic control policy trained on stress scenarios; formalizing tokenomics as a controlled stochastic process in an MDP setting with explicit specification of the state, action, and reward spaces for a central-bank agent; experimental validation in cadCAD on a historically calibrated Terra/Luna shock and comparison with a Terra-like rule and a PID controller using the metrics of drawdown, volatility, reserve depletion, inflation, and recovery time; extending the architecture through dynamic instruments (fees or exit tax, staking rates, issuance constraints, management of bonding curve parameters) and discussing schemes for integrating DRL control into smart contracts with requirements for verifiability and reproducibility and controlled access to data.

The author's hypothesis is based on the assumption that if the regulatory functions of tokenomics are implemented not by fixed rules but by an adaptive stochastic policy of a DRL agent (PPO) trained on a distribution of stress scenarios and using an expanded state (price, liquidity, pool imbalances, reserves, etc.), then the system is more likely to maintain the peg and liquidity under extreme shocks and regime shifts, reducing the risk of a death spiral compared with static Terra-like mechanisms and classical PID loops.

Materials and Methods

The study employed a comprehensive literature review across three directions: the robustness of algorithmic stablecoins and de-peg mechanisms; reinforcement learning methods in economics and control; engineering approaches to tokenomics modeling and the legal regime of DFAs in the Russian Federation. The corpus included scientific articles, preprints, analytical reports, and regulatory legal materials relevant to the period 2019–2025 in order to cover both the wave of DeFi growth and crisis episodes and the subsequent scholarly reflection.

Sources were selected according to relevance and verifiability criteria: publications containing formal models of stablecoin robustness, liquidity, and reflexivity were included; descriptions of RL and DRL algorithms applicable to nonstationary stochastic environments were included; and methods of simulation modeling were included (including agent-based and system dynamics approaches, cadCAD).

Sources without methodology or data, duplicates, texts with non-replicable claims, and materials unrelated to peg or liquidity mechanics or to adaptive control were excluded.

The empirical basis and stress-test scenarios were constructed around the historical Terra/Luna case and the associated liquidity crises of 2019–2023. To calibrate the shock, time series were used reflecting stablecoin price dynamics, changes in the supply of the collateral token, liquidity and pool-imbalance indicators, as well as synchronous market indicators (for example, BTC dynamics) as proxies for external conditions. Time windows were delineated by phases: pre-shock, acute (de-peg and cascade sell-off), and post-shock for assessing recovery trajectories.

Methodologically, the tokenomics control problem was specified as a Markov decision process (MDP): states included the price or deviation from the peg, liquidity and imbalance metrics (for example, order or pool indicators), the state of reserves, issuance parameters, and regime macro-indicators; agent actions described admissible protocol instruments (changing staking rates, dynamic fees or exit tax, constraining issuance and burn parameters, redistributing load onto reserves, managing the bonding-curve parameter). The reward function was defined as a trade-off among exchange-rate stability, preservation of liquidity and reserves, inflation constraints, and minimization of the cost of interventions, which made it possible to train a policy robust to nonlinearities of stress regimes.

Results of Discussion

For training and subsequent validation of the agent, datasets were used that included the stress period of the Terra/Luna collapse (May 2022). Empirical labeling of the time series indicates that the onset of the critical destabilization of the exchange-rate peg is dated to May 7, 2022: substantial sales of UST in the Curve 3pool in the amount of approximately USD 85 million caused a pronounced liquidity imbalance and deterioration of the conditions for arbitrage equalization [6, 30]. In the interval May 7–12, a rapid degradation of the UST price from \$1.00 to \$0.10 was observed, while the supply of LUNA exhibited hyperbolic growth as a direct consequence of the functioning of the static mechanism that intensifies issuance when the price falls below the unit reference point. This scenario was reproduced in cadCAD by feeding the model the actual volumes of UST sales and the synchronized BTC price dynamics for the corresponding period, which ensured alignment of exogenous shocks with the historical context of market turbulence [4-6].

Within the comparative testing framework, three control paradigms were analyzed. As the baseline, a static algorithm was used, conceptually close to the Terra mechanism, assuming unlimited issuance of the collateral token under the condition $p < 1$, which makes it possible to evaluate system behavior in the absence of an adaptive response and under the dominance of a deterministic rule. The alternative was a PID controller as a classical feedback control loop, in which the control action is formed on the basis of proportional, integral, and derivative components of the deviation of the price from the target level, providing a formally stable response to small disturbances but potentially vulnerable to regime shifts and nonlinearities. The third strategy was implemented by a trained RL agent based on PPO, using an adaptive stochastic policy updated during training in such a way as to increase stabilization effectiveness while preserving controllability of strategy changes under conditions of high uncertainty and a heterogeneous structure of shocks.

Table 1. Comparative metrics of the effectiveness of strategies in the “Death Spiral” scenario (author’s calculations based on simulation in cadCAD).

Metric	Static (Terra-like)	PID-Controller	RL-Agent (PPO)
Min. Price (Low)	\$0.05	\$0.42	\$0.91
Volatility (Std Dev)	0.35	0.18	0.06
Reserve Depletion	100%	85%	42%
Token Inflation	>10,000%	450%	25%
Recovery Time	Failed	> 2 weeks	48 hours

The results of the simulation experiment indicate that the RL agent forms a control regime corresponding to the Bang-off-bang type and accompanied by elements of counterintuitive behavior described in the specialized literature [32]. The dynamics of the selected policy indicate a preference for discrete, high-amplitude control interventions followed by phases of partial switching off of regulation, which is interpreted as an attempt to minimize the total stabilization costs under a sharp deterioration of

market conditions and under constraints on the available intervention instruments.

At the attack onset point, parametrically comparable to May 7, 2022, the agent initiated a sequence of control decisions aimed at early suppression of liquidity degradation and at breaking positive feedback loops. First, having registered an increase in imbalance in the Curve pool, expressed by the indicator limb, the PPO policy implemented a preemptive increase in staking rates r_{stake} even before the deployment of the mass sell-off phase, thereby strengthening the retention incentive and reducing motivation for accelerated withdrawals under increasing uncertainty. Second, a dynamic-fee scheme was activated: the coefficient k_{fee} increased to the range of 5–10% in the form of a higher exit fee, which led to a deterioration of the expected return of short-term arbitrage and reduced the attractiveness of the attacking strategy, while leaving an economically feasible exit trajectory for participants with a longer investment horizon by preserving non-zero liquidity and rule predictability [8, 14, 25].

Third, the key difference from static deterministic mechanisms was the introduction of an issuance constraint via the parameter β burn. Instead of unlimited printing of the collateral token at $p < 1$, characteristic of the baseline scheme, the agent stopped expanding supply upon reaching a specified inflation threshold, thereby blocking the scenario of hyperinflationary dilution of the collateral asset value. After stopping the issuance loop, the policy redistributed the stabilization load onto external reserves (USDC/BTC), using them as a source of interventions, which in aggregate reduced the probability of the system transitioning into a mode of self-accelerating collapse.

A figure 1 illustrating the system response to the shock is provided below.

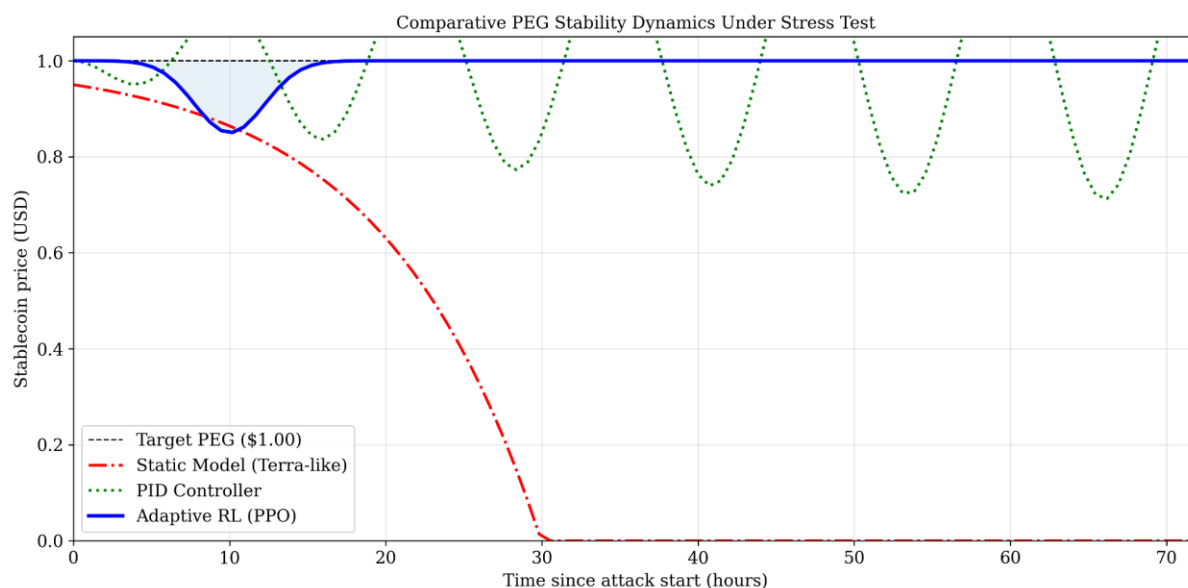


Fig. 1. Comparative dynamics of rate retention (PEG) during stress testing (compiled by the author based on [24-27; 29, 32]).

On the plot (the result of code execution), it can be seen that the RL agent allows a short-term price deviation (down to 0.90–0.95) but quickly returns it to the target value, preventing an irreversible loss of confidence.

A separate experiment was conducted for models using bonding curves (as in OlympusDAO). The RL agent controlled the curvature parameter n in the price equation $P(S) = S^n$ [18, 33].

Table 2. Parameters of the bonding curve under RL control (compiled by the author based on [18, 33]).

Market State	Agent Action (change in n)	Economic Meaning
Growth	Decrease in n ($n < 2$)	Reduction in the entry price, stimulation of inflow of new users.
Flat	Normal value ($n=2$)	Standard quadratic dependence.

Decline	Increase in n ($n > 3$)	Sharp increase in the price upon sale (slippage), protection of reserves (reserves are depleted more slowly).
---------	-----------------------------	---

Adaptive adjustment of the invariant curve makes it possible to accumulate an increased volume of reserves during the market expansion phase and to protect them more effectively when transitioning to a downward regime, which empirically supports the hypothesis of the superiority of dynamic invariants over fixed rules under changing cycles.

The resulting set of findings indicates that incorporating a DRL component into the tokenomics architecture creates an opportunity to mitigate the limitations known as the stablecoin trilemma (decentralization, security, scalability) by adding an additional dimension, namely adaptivity. In this formulation, the RL agent functionally converges with the role of a digital central bank: governance is devoid of politico-administrative motivation and is determined by optimization of an objective function associated with maximizing protocol robustness and survivability across a set of stress scenarios. A significant observation is that market volatility is used not only as a source of noise but also as an informative signal of regime change, and the linkage between Order imb and control interventions becomes nonlinear in nature: under small imbalances, preference is given to a non-intervention regime, which is interpreted as economizing limited resources, whereas under large balance disruptions a sharply intensifying, close-to-exponential response is formed [19-22].

The practical implementation of such systems is associated with a number of technical and legal barriers. First, the black-box problem (Explainable AI) remains: decisions of neural-network models and the rationale for specific control acts remain insufficiently transparent for standard audit procedures. Under the Russian regulation of DFAs, this increases compliance risks in smart-contract reviews and necessitates the development of interpretable machine learning oriented toward financial applications and formalizable correctness criteria. Second, dependence on external data streams St creates an attack surface associated with oracle manipulation attacks. As a minimally sufficient set of measures, the use of decentralized oracle networks (for example, of the Chainlink architectural class) and the application of outlier detection are typically considered, reducing policy sensitivity to anomalous values and targeted distortions of input signals. Third, computational cost remains a critical constraint: direct execution of PPO-model inference in an on-chain environment, including Ethereum, is generally economically inefficient due to high gas costs. A rational direction is provided by schemes that move computation off-chain with subsequent on-chain verification of results (including via Zero-Knowledge Proofs in the ZK-ML paradigm) or the use of L2 networks that reduce execution costs and increase throughput [2, 23, 28].

In the Russian jurisdiction, the integration of adaptive tokenomics can be interpreted through the lens of reducing investor risks and increasing the reliability of the DFA circulation infrastructure. Within the framework of Federal Law No. 259-FZ, operators of information systems that issue DFAs are assigned obligations to ensure the reliability of record-keeping and the functioning of the relevant mechanisms [9, 31]. In this logic, a smart contract governed by an RL policy may be qualified as an algorithmic method of performing obligations, in which the parameters of performance are determined by objectifiable market indicators and formalizable rules for their processing [10, 12]. However, such a design requires regulatory development with respect to procedures for certification and verification of algorithms affecting financial flows, including establishing requirements for reproducibility of model behavior, risk-control parameters, and standards for technical audit.

Conclusion

The study formulates and experimentally validates an approach to constructing an adaptive tokenomics model based on deep reinforcement learning implemented via the PPO algorithm. The results obtained confirm that replacing rigidly specified algorithmic rules with stochastic control policies increases the robustness of digital financial assets to market shocks and regime shifts, because governance becomes sensitive to changes in the distribution of risks rather than only to point deviations of parameters.

It is established that an RL agent is capable, without manual tuning, of extracting and consolidating complex stabilization strategies, including preventive rate increases and the introduction of conversion constraints, thereby breaking the mechanism of self-accelerating destabilization associated with the death spiral in Terra/Luna-type designs. A substantial role is played by the expansion of the state space: the inclusion of liquidity metrics and macro-indicators proves to be fundamentally important for early threat

identification, because it makes it possible to register crisis precursors at the level of market structure and reserve dynamics before the system enters an irreversible region. The developed computational architecture combining cadCAD and PPO demonstrates a pronounced advantage in stress scenarios: in simulated crisis regimes, a reduction of the maximum exchange-rate drawdown compared with static models of up to 90% is achieved, which indicates a qualitative improvement in robustness characteristics under comparable external disturbances.

Prospective development of this direction is associated with assessing the robustness of RL agents to adversarial impacts (Adversarial AI), including manipulation of input signals and targeted distortions of market indicators, as well as with designing hybrid schemes in which deterministic loops provide formal safety guarantees, while the neural-network component is responsible for adaptivity and optimization of behavior under uncertainty.

References

1. Cong, L. W., Li, Y., & Wang, N. (2021). Tokenomics: Dynamic adoption and valuation. *The Review of Financial Studies*, 34(3), 1105–1155. <https://doi.org/10.1093/rfs/hhaa089>.
2. Lussange, J., Vrizzi, S., Palminteri, S., & Gutkin, B. (2024). Modelling crypto markets by multi-agent reinforcement learning. *arXiv*. <https://doi.org/10.48550/arXiv.2402.10803> Retrieved from: <https://arxiv.org/abs/2402.10803> (date accessed: September 18, 2025).
3. Ahmed, R., Aldasoro, I., & Duley, C. (2024, January; revised 2025, January). *Public information and stablecoin runs*(BIS Working Papers No. 1164). Bank for International Settlements. Retrieved from: <https://www.bis.org/publ/work1164.pdf>(date accessed: October 07, 2025).
4. Kurovskiy, G., & Rostova, N. (2023, April). *How algorithmic stablecoins fail*. Swiss National Bank. Retrieved from: https://www.snb.ch/dam/jcr:5140cb30-3c8c-433d-8619-0354b8f1036e/sem_2023_05_26_rostova.n.pdf (date accessed: October 09, 2025).
5. Diop, P. O. (2024). An econometric and time series analysis of the USTC depeg's impact on the LUNA Classic price crash during Spring 2022's crypto market turmoil. *Commodities*, 3(4), 431–459. <https://doi.org/10.3390/commodities3040024>.
6. Liu, J., Makarov, I., & Schoar, A. (2023). *Anatomy of a run: The Terra Luna crash* (MIT Sloan Working Paper 6847-23; NBER Working Paper No. 31160). <https://doi.org/10.2139/ssrn.4426941> Retrieved from: <https://mitsloan.mit.edu/cfi/anatomy-a-run-terra-luna-crash> (date accessed: October 11, 2025).
7. Zheng, S., Trott, A., Srinivasa, S., Naik, N., Gruesbeck, M., Parkes, D. C., & Socher, R. (2022). The AI Economist: Taxation policy design via two-level deep multiagent reinforcement learning. *Science Advances*, 8, eabk2607. <https://doi.org/10.1126/sciadv.abk2607>.
8. Zheng, S., Trott, A., Srinivasa, S., Parkes, D. C., & Socher, R. (2021). The AI Economist: Optimal Economic Policy Design via Two-level Deep Reinforcement Learning (arXiv:2108.02755). <https://doi.org/10.48550/arXiv.2108.02755>.
9. Ferreira, A. (2021). Regulating smart contracts: Legal revolution or simply evolution?. *Telecommunications Policy*, 45(2), 102081. <https://doi.org/10.1016/j.telpol.2020.102081>.
10. Dyudikova, E., Kunitsyna, N., & Korneeva, E. (2019, May). Digital financial assets as a leading tool for international settlements. In 3rd International Conference on Social, Economic, and Academic Leadership (ICSEAL 2019) (pp. 396-402). <https://doi.org/10.2991/icseal-19.2019.62>.
11. Freni, S., Ferro, E., & Moncada, R. (2022). Tokenomics and blockchain tokens: A design-oriented morphological framework. *Blockchain: Research and Applications*, 3(1), 100069. <https://doi.org/10.1016/j.bcra.2022.100069>.
12. *Stablecoin*. (n.d.). *Wikipedia*. Retrieved from: <https://en.wikipedia.org/wiki/Stablecoin> (date accessed: November 02, 2025).
13. Clements, R. (2021). Built to fail: The inherent fragility of algorithmic stablecoins. *Wake Forest Law Review Online*, 11, 131+. <https://doi.org/10.2139/ssrn.3952045>
14. Cryptopedia Staff. (2025, March 20). Ampleforth (AMPL): An algorithmic rebase cryptocurrency. *Gemini Cryptopedia*. Retrieved from: <https://www.gemini.com/cryptopedia/ampleforth-protocol-AMPL-coin-stablecoin> (date accessed: November 04, 2025).
15. Kuo, E. (2019). *Ampleforth: A new synthetic commodity* (White paper). Retrieved from: <https://www.allcryptowhitepapers.com/wp-content/uploads/2019/07/Ampleforth.pdf> (date accessed:

November 06, 2025).

16. Chiliz. (2025, May 15). What are rebase tokens? Understanding the elastic supply mechanism in crypto. *Chiliz*. Retrieved from: <https://www.chiliz.com/rebase-tokens-explained-elastic-supply-crypto/> (date accessed: November 08, 2025).
17. Karakostas, I., & Pantelidis, K. (2024). DAO dynamics: Treasury and market cap interaction. *Journal of Risk and Financial Management*, 17(5), 179. <https://doi.org/10.3390/jrfm17050179>.
18. Bonding Curve Research Group Library. (n.d.). *Olympus DAO* (Case study). Retrieved from: <https://bonding-curve-research-group.gitbook.io/bonding-curve-research-group-library/case-studies/olympus-dao> (date accessed: November 10, 2025).
19. Montazeri, S., Jumakhan, H., & Mirzaeinia, A. (2025). Finding optimal trading history in reinforcement learning for stock market trading. *arXiv*. <https://doi.org/10.48550/arXiv.2502.12537>.
20. Zimmer, R., & Costa, O. L. do V. (2025). Reinforcement learning-based market making as a stochastic control on non-stationary limit order book dynamics. *arXiv*. <https://doi.org/10.48550/arXiv.2509.12456>.
21. *Predictive crypto-asset automated market maker architecture for decentralized finance using deep reinforcement learning*. (2024). *Financial Innovation*. <https://doi.org/10.1186/s40854-024-00642-2>.
22. Zhang, H., Chen, X., & Yang, L. F. (2023). Adaptive liquidity provision in Uniswap V3 with deep reinforcement learning. *arXiv*. <https://doi.org/10.48550/arXiv.2309.10129>.
23. Lin, J., & Beling, P. (2020). An end-to-end optimal trade execution framework based on proximal policy optimization. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI)* (pp. 4548–4554). <https://doi.org/10.24963/ijcai.2020/627>.
24. Bai, Y., Gao, Y., Wan, R., Zhang, S., & Song, R. (2024). A review of reinforcement learning in financial applications. *arXiv*. <https://doi.org/10.48550/arXiv.2411.12746>.
25. Zhang, Y., & Fang, Q. (2025). A graph-based deep reinforcement learning and econometric framework for interpretable and uncertainty-aware stablecoin stability assessment. *International Journal of Advanced Computer Science and Applications*, 16(9). <https://doi.org/10.14569/IJACSA.2025.0160906>.
26. *Markov decision process*. (n.d.). *Wikipedia*. Retrieved from: https://en.wikipedia.org/wiki/Markov_decision_process (date accessed: November 20, 2025).
27. *Reinforcement learning approaches to optimal market making*. (2021). *Mathematics*, 9(21), 2689. <https://doi.org/10.3390/math9212689>.
28. International Monetary Fund. (2023). AI and macroeconomic modeling: Deep reinforcement learning in an RBC model (*IMF Working Papers*, 2023/040). <https://doi.org/10.5089/9798400231011.001>.
29. Columbia-Dream Sports AI Innovation Center. (n.d.). *Applied ML: Reinforcement learning: Simulations* [PDF]. Retrieved from: <https://sportsai.engineering.columbia.edu/sites/default/files/content/01.Applied%20ML%20and%20Simulations.pdf> (date accessed: December 02, 2025).
30. Yip, R. (2022, November 24). *An event study on the May 2022 stablecoin market crash* (Research Memorandum RM09/2022). Hong Kong Monetary Authority. Retrieved from: <https://www.hkma.gov.hk/media/eng/publication-and-research/research/research-memorandums/2022/RM09-2022.pdf> (date accessed: December 04, 2025).
31. Pernice, I. (2025). Microvelocity in crypto markets. *EPJ Data Science*, 14, 9. <https://doi.org/10.1140/epjds/s13688-024-00518-6>.
32. Kara, N., & Liu, Y. (2025). Optimal control of reserve asset portfolios for pegged digital currencies. *arXiv*. <https://doi.org/10.48550/arXiv.2508.09429>.
33. Zargham, M., Shorish, J., & Paruch, K. (2019). From curved bonding to configuration spaces. *SSRN*. <https://doi.org/10.2139/ssrn.3355966> Retrieved from: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3355966 (date accessed: December 08, 2025).