Proceedings of the 5th International Conference on Financial Technology and Business Analysis
DOI: 10.54254/2754-1169/2026.38909

Does Economic Momentum Exist in VALORANT? A Discrete-Time Markov Chain Analysis with Logistic Regression and Random Forest Validation

Yaoming Yu

A-Level International Curriculum Center, Harbin Engineering University, Harbin, China
13804512621@139.COM

Abstract. Does winning beget more winning in VALORANT? This paper examines whether teams that pull ahead financially keep pulling ahead, or whether the game's economy naturally evens things out. A four-state Markov chain is fitted to transitions between economic regimes across 214 professional matches (5,136 rounds) from the 2023–2024 VCT season. Logistic regression and a random forest model are also used to test whether the patterns hold up across different modeling approaches. The results are hard to ignore: winners stay in full-buy 78% of the time, while losers remain there only 31%. Each extra 1,000 credits of advantage raises the odds of winning by about one-third (OR = 1.34, $p < 0.001$). Adding regime labels increases AUC from 0.68 to 0.72, representing a 5.9% gain. Overall, the findings suggest that money begets money in VALORANT and that the in-game economy has memory. Economic regimes provide useful information beyond static credit numbers alone and help explain how financial momentum develops across rounds.

Keywords: Economic momentum, VALORANT, Markov chain, logistic regression, esports analytics

1. Introduction

Tactical first-person shooters like VALORANT are built around an economy. Every round, teams earn credits based on how they did and spend them on guns, shields, and abilities. Win a round, and next round comes easier. Lose, and climbing back can take several rounds. The question this paper asks is simple: does that advantage just keep compounding, or does the game pull teams back to balance? This is not just idle curiosity. If economic momentum is real, it changes how teams should plan buys across a half, and how analysts should read a match as it unfolds.

Markov chains have been used in sports for decades. Tennis was one of the first targets [1], and the same toolkit later showed up in basketball [2], and baseball salary analysis [3]. In esports, most of the existing work goes in the opposite direction: rather than asking why patterns happen, people build models to predict who wins. Hodge et al. ran logistic regression on multi-player esports matches [4]; Yang et al. built a real-time gradient-boosted system [5]. More recently, Birant and Birant pushed accuracy to 95% with a multi-instance approach [6], and Chowdhury and Ahsan benchmarked linear and tree-based models on League of Legends [7]. Xenopoulos et al. found that

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2754-1169/2026.38909

resource gaps matter more at higher skill levels [8]. StarCraft II research points the same way [9], and Counter-Strike economic decisions have been modeled directly [10]. VALORANT specifically has had economy-based predictors studied too [11, 12], alongside broader FPS modeling work [13].

Two things are missing from this picture. One, most papers optimize for prediction accuracy and treat each round as a fresh event, ignoring that rounds are strung together. Two, Markov models are common in traditional sports but have barely been applied to esports economies. Nobody has really put the question directly: does economic advantage snowball, or is it just a matter of who has more credits right now?

This paper takes the snowball question seriously. Rather than chasing a better predictor, the goal is to figure out whether momentum is actually there. A four-state Markov chain labels each team's economy as full-buy, half-buy, eco, or force-buy. Logistic regression then checks whether those states and the credit gap predict winning, and a random forest runs the same job to make sure the answer is not an artifact of one model. The data come from 214 VCT matches, about 5,000 rounds in total. Three questions are on the table: do winners and losers move between states differently? Does a credit gap actually predict the next round? Does knowing the current regime add anything on top of just knowing the credit numbers?

2. Methodology

2.1. Data source and preprocessing

Data comes from VLR.GG, a community-maintained statistics site for professional VALORANT. The sample covers 214 matches across the 2023-2024 VCT season, giving 5,136 rounds. Matches span every major region - Americas, EMEA, Pacific, China - plus international events like Masters and Champions. Lower-tier and exhibition matches were filtered out to keep the competition level consistent.

Per round, the following were recorded: map, round number, starting credits per player, equipment value, who won, and which side (attack or defense). Credit differential is just the gap in total starting credits between the two teams, scaled to thousands. Equipment gap is built the same way from equipment value. Regimes are assigned by the standard community split: 3,900+ credits per player is full-buy; 2,000 to 3,899 is half-buy; under 2,000 gets split by spending rate - spending 80% or more of available credits is force-buy, otherwise eco.

2.2. Model construction

The Markov chain has four states, one per regime. The Markov assumption here is that the next state depends only on the current one. Transition probabilities are just the raw counts: from state i to state j, divide the number of times i turned into j by how many times i showed up at all. The long-run distribution comes from solving pi = pi P, which tells what fraction of rounds land in each regime in equilibrium.

Logistic regression tests whether the credit gap actually predicts wins. The outcome is binary: win or lose. The model takes credit differential, equipment gap, and map side as predictors. Coefficients are read as odds ratios - multiply by e to get the multiplicative change in win odds per unit of predictor.

A random forest is also run as a cross-check. Trees are known for picking up non-linear effects and interactions that a straight logistic regression will miss. Here 500 trees were grown, splitting on

DOI: 10.54254/2754-1169/2026.38909

Gini impurity. The metric for both models is AUC, the standard measure of how well a model ranks wins above losses.

2.3. Evaluation metrics

The checks are straightforward. Transition asymmetry gets a chi-squared test comparing full-buy retention between winners and losers. Model comparison uses DeLong's test for correlated ROC curves. Diagnostics cover VIF (multicollinearity), Hosmer-Lemeshow (fit), and 10-fold cross-validation (stability). McFadden's pseudo R-squared is reported for the regression model.

3. Results

3.1. Descriptive statistics

Regime counts are in Table 1. Full-buy shows up in 41% of rounds, more than any other state. Half-buy and eco are close behind at 23% and 22%, with force-buy at 14%. The average credit gap is 1,200 credits, but that average hides a lot - the standard deviation is about 2,500, so blowout gaps are common. The team with more starting credits wins 58.4% of the time, which is already a hint that money matters (binomial test, $p < 0.001$).

Table 1. Distribution of economic regimes

Economic Regime	Frequency	Percentage
Full-Buy	2,106	41.0%
Half-Buy	1,181	23.0%
Eco	1,130	22.0%
Force-Buy	719	14.0%

3.2. Hypothesis 1: transition asymmetry

The transition matrix is Table 2. The diagonal is what matters most - it tells how likely a team is to stay in the same regime. Full-buy has the stickiest diagonal at 0.72: once there, teams tend to stay. Eco is much more slippery, with self-transition at only 0.23. From eco, teams either go all-in to force-buy (0.41) or bounce back to full-buy (0.36). One structural feature stands out: nobody in full-buy or half-buy goes directly to force-buy, which makes sense - force-buy is by definition about spending almost everything when you are already broke.

Table 2. Estimated transition probability matrix (MLE)

From / To	Full-Buy	Half-Buy	Eco	Force-Buy
Full-Buy	0.72	0.28	0.00	0.00
Half-Buy	0.15	0.30	0.55	0.00
Eco	0.36	0.00	0.23	0.41
Force-Buy	0.62	0.38	0.00	0.00

DOI: 10.54254/2754-1169/2026.38909

The long-run distribution is [0.412, 0.228, 0.217, 0.143], confirming full-buy as the default state. Now the real test: splitting the matrix by who won the round. Winners stay in full-buy 78% of the time; losers only 31%. That gap is huge and significant (chi-sq = 247.3, df = 1, $p < 0.001$). A goodness-of-fit test does not reject the first-order assumption (chi-sq = 18.7, df = 15, p = 0.23), so a one-step memory looks like a reasonable approximation.

3.3. Hypothesis 2: resource inequality and outcome

Table 3 has the regression output. Credit differential is the dominant predictor: OR = 1.34 per 1,000 credits (95% CI: 1.21-1.49, $p < 0.001$). In plain terms, every extra 1,000 credits raises win odds by about a third. A 3,000-credit gap - roughly the difference between full-buy and eco - multiplies the odds by 2.4. Equipment gap matters too (OR = 1.18, p = 0.003) but less than raw credits. Attacking side shows a slight negative effect (OR = 0.92) that almost but not quite reaches significance (p = 0.087).

Table 3. Logistic regression results

Variable	Odds Ratio	95% Confidence Interval	p-value
Credit Differential	1.34	[1.21, 1.49]	< 0.001 ***
Equipment Gap	1.18	[1.07, 1.30]	0.003 **
Map Side (Attack)	0.92	[0.84, 1.01]	0.087 n.s.

The baseline regression scores 0.68 AUC - not a crystal ball, but well above the 0.5 baseline and in the normal range for round-level esports prediction. McFadden pseudo R-squared is 0.09, which is typical when so much of the outcome boils down to raw team skill and luck. The random forest lands at 0.69 AUC, essentially the same story. Its feature importance ranking lines up with the regression: credit differential on top, then equipment gap, then side.

3.4. Hypothesis 3: incremental value of temporal features

Adding regime dummies to both models tests whether knowing the current economic state helps beyond the credit numbers. For logistic regression, AUC moves from 0.68 to 0.72 - a 5.9% relative gain. The DeLong test says this is real (z = 2.41, p = 0.016), and the likelihood ratio test agrees (chi-sq = 18.3, df = 6, p = 0.012). The random forest barely moves (0.69 to 0.71), probably because trees already pick up some of the regime effects through interactions.

Digging into the regime coefficients, force-buy and eco both predict lower win probability even after controlling for how much money is on the table. That suggests the state itself carries a penalty beyond the resource gap - maybe confidence, maybe playstyle, maybe both. Diagnostics look fine: VIF all under 3, Hosmer-Lemeshow p-values at 0.21 and 0.18, and cross-validation stable (SD = 0.023).

4. Discussion

The 78% versus 31% gap is the headline. Winning teams do not just happen to be rich; they stay rich. Losing teams, by contrast, cycle out of full-buy fast. That is momentum in plain language: a positive feedback loop where winning funds the next win. The first-order assumption holding up adds weight to the idea that the economy has memory, not just random noise round to round.

DOI: 10.54254/2754-1169/2026.38909

One detail that deserves attention is that the credit gap does not hit equally everywhere. In half-buy rounds, small equipment differences swing outcomes a lot. When both sides are fully kitted out, an extra rifle changes little. This explains why regime labels help prediction: they tell you how much a given credit gap actually matters. The fact that eco and force-buy teams underperform even after matching on credits hints at something beyond gear itself - tilt, cautious play, or bad decision-making when money is tight.

The 5.9% AUC gain from adding regimes is real but not dramatic. This matches what sports analytics has always found: contextual features squeeze out incremental gains, while the bulk of predictive power sits in the basic numbers. In a betting or coaching setting, even a few points can matter over hundreds of predictions, but for understanding the game, static indicators still carry most of the signal.

Getting the same answer from a linear model and a tree ensemble helps. If only logistic regression had shown these effects, one might wonder whether specification drove the result. Both models land on credit differential as the top feature with nearly identical AUCs, which is a decent sign that the relationship is there regardless of how it is modeled.

5. Conclusion

The question was whether economic momentum exists in VALORANT. It does. Winners and losers move through economic regimes very differently (78% versus 31% full-buy retention), a credit gap of 1,000 credits moves win odds by about a third (OR = 1.34, $p < 0.001$), and knowing the current regime adds about 5.9% to AUC on top of static credit numbers. In short, the economy remembers where it has been. Three caveats apply. The first-order Markov assumption may miss second-order psychological effects, the sample is pro-only so lower ranks might look different, and variables like team chemistry and agent composition are not in the model. Future work could stretch across patches, try second-order models, include lower-skill data, and apply the same framework to other games like CS2.

References

[1] Newton, P. K., & Keller, J. B. (2005). Probability of winning at tennis I: Theory and data. *Studies in Applied Mathematics, 114*(3), 241–269. https://doi.org/10.1111/j.0022-2526.2005.01539.x

[2] Grinstead, C. M., & Snell, J. L. (2012). *Introduction to probability* (2nd ed.). American Mathematical Society. https://www.dartmouth.edu/~chance/teaching_aids/books_articles/probability_book/book.html

[3] Annala, C. N., & Winfree, J. (2011). Salary distribution and team performance in Major League Baseball. *Sport Management Review, 14*(2), 167–175. https://doi.org/10.1016/j.smr.2010.08.002

[4] Hodge, V. J., et al. (2019). Win prediction in multi-player esports: Live professional match prediction. *IEEE Transactions on Games, 13*(1), 68–79. https://doi.org/10.1109/TG.2019.2955562

[5] Yang, Y., Qin, T., & Lei, Y.-H. (2016). Real-time eSports match result prediction. arXiv preprint arXiv:1701.03162. https://arxiv.org/abs/1701.03162

[6] Birant, K. U., & Birant, D. (2023). Multi-objective multi-instance learning: A new approach to machine learning for eSports. *Entropy, 25*(1), 28. https://doi.org/10.3390/e25010028

[7] Chowdhury, S., & Ahsan, M. (2025). Applications of linear and ensemble-based machine learning for predicting winning teams in League of Legends. *International Journal of Computer Science and Network Security, 25*(3), 112–121.

[8] Xenopoulos, P., Freeman, W. R., & Silva, C. (2022). Analyzing the differences between professional and amateur esports through win probability. In *Proceedings of the ACM Web Conference 2022* (pp. 1763–1773). ACM. https://doi.org/10.1145/3485447.3512218, New York.

[9] Białecki, A., et al. (2022). Determinants of victory in esports: StarCraft II. *Multimedia Tools and Applications, 81*(28), 40157–40172. https://doi.org/10.1007/s11042-022-12572-9

DOI: 10.54254/2754-1169/2026.38909

[10] Xenopoulos, P., Coelho, B., & Silva, C. (2021). Optimal team economic decisions in Counter-Strike. *arXiv preprint arXiv:2109.12990*. https://arxiv.org/abs/2109.12990
[11] Pawar, A. M. (2024). *Valorant esports pre-match betting advisory system* (MSc thesis). National College of Ireland.
[12] Wang, D. (2024). A predictive analysis of Valorant esports: Win probability through economy and ultimate ability. *TechRxiv*. DOI: 10.36227/techrxiv.174742078. Available at: https://www.techrxiv.org/articles/preprint/A_Predictive_Analysis_of_Valorant_Esports_Win_Probability_Through_Economy_and_Ultimate_Ability/174742078
[13] Xia, X., Salinas, A., & Morstatter, F. (2025). Precision under fire: Analysis and predictive modeling in Counter-Strike: Global Offensive. In *Proceedings of the IEEE Conference on Games (CoG)*. DOI: 10.1109/CoG64752.2025.11114164.