

# ***Casino-Side Kelly Analysis of Roulette: Growth, Fractional Exposure, and Finite-Horizon Risk***

**Chenxin Weng Wang**

*Eötvös József Gimnázium, Budapest, Hungary*  
*wengchenxin2009@gmail.com*

**Abstract.** Roulette gives a transparent setting for studying how a small statistical advantage can be converted into a controlled level of repeated financial exposure. This study develops a casino-side analysis based on the Kelly criterion, comparing American and European even-money roulette under full, half, and quarter Kelly strategies. The theoretical analysis derives the growth-maximising exposure from expected logarithmic wealth and uses concavity to establish uniqueness. The model is then calibrated to the two wheel structures and evaluated through Monte Carlo simulation over 1,000 spins. For American roulette, full Kelly produces the highest median terminal bankroll, but also a substantially weaker lower tail and a greater probability of finishing below the initial bankroll than fractional Kelly. European roulette exhibits the same pattern at a smaller scale because its house advantage is lower. These results show that long-run optimality does not remove finite-horizon risk: a strategy can maximise compound growth while still generating deep temporary losses. Fractional Kelly therefore provides a practical compromise when capital constraints, operating expenses, or model uncertainty make downside control important. The roulette benchmark also clarifies how the framework can be extended to blackjack, sports odds, and portfolio allocation, where probabilities and payoffs must be estimated rather than read directly from fixed game rules.

**Keywords:** Kelly criterion, roulette, fractional betting, logarithmic growth, Monte Carlo simulation

## **1. Introduction**

Roulette presents a simple but economically meaningful risk-allocation problem. A casino repeatedly accepts wagers with fixed payoffs, yet the outcome of each spin remains uncertain. American roulette has 38 pockets, including 0 and 00, whereas European roulette has 37 pockets with only 0. On an even-money red or black bet, these additional pockets create a known structural advantage for the casino [1]. The operational problem is therefore not whether an advantage exists, but how strongly the casino should expose its bankroll to that advantage when losses can occur in long sequences.

The Kelly criterion provides a natural framework for this problem because it converts a probabilistic edge into a proportional exposure that maximises expected logarithmic wealth. Kelly connected this objective to long-run capital growth [2]. Breiman established central asymptotic

properties of favourable gambling systems [3], and Thorp developed the framework for practical gambling applications, including blackjack [4]. A later collection edited by MacLean, Thorp, and Ziemba systematised the theory and explained why fractional Kelly strategies can reduce short-run risk at the cost of some long-run growth [5].

Recent research has shifted attention toward constraints, estimation error, and finite horizons. Reinforcement-learning work has treated estimation error as a practical weakness of full Kelly and compared it with fractional or shrinkage alternatives [6]. Fractional Kelly has also been extended to non-mutually exclusive bets [7], while portfolio research has related growth-optimal Kelly allocations to risk-adjusted efficient strategies [8]. For repeated games, the optimal fraction may depend on the remaining horizon and a maximum payout [9]. Distributional analysis has further shown that a Kelly-style model can optimise different bankroll quantiles and represent ruin risk explicitly [10]. Data-driven robust formulations similarly address uncertainty in the return distribution and seek improved stability and drawdown control [11].

Against this background, the present study uses roulette as an exact-probability benchmark. Its objective is to measure how a known house edge translates into a growth-optimal casino exposure and how much long-run growth is surrendered when exposure is reduced to control finite-horizon downside. The analysis derives the general binary Kelly rule, calibrates it to American and European roulette, and compares full, half, and quarter Kelly through Monte Carlo simulation. The contribution is a casino-side interpretation that separates the source of the edge from the choice of exposure and connects the theoretical optimum to operational risk.

## 2. Methods

### 2.1. General Kelly exposure

Let  $W$  denote the casino bankroll before a wager and let  $f$  denote the fraction exposed, where the fraction is non-negative and less than one. The casino wins the stake with probability  $p$  and loses with probability  $q$ , with the two probabilities summing to one. If a casino win earns net odds  $b$  per unit exposed, the next bankroll is  $W(1 + bf)$  after a win and  $W(1 - f)$  after a loss. The arithmetic expected return per unit exposed is  $bp - q$ , so the wager is favourable to the casino only when that quantity is positive.

Logarithms are appropriate because repeated wealth factors multiply while their logarithms add. Under independent repetition, the sample average log return converges to a compound-growth rate, and outcomes close to ruin receive a severe penalty [2, 3].

Strict concavity makes any interior critical point the unique maximum. The criterion does not create an advantage: positive exposure appears only when the underlying expected payoff is positive. If the unconstrained solution is negative, the simplified model prescribes zero exposure; if it exceeds an operational limit, the constrained optimum occurs at that limit.

### 2.2. Roulette calibration and fractional exposure

For an even-money roulette bet, the net odds equal one. From the casino perspective, the win probabilities are  $20/38$  for an American wheel and  $19/37$  for a European wheel. The resulting full-Kelly exposures equal the respective house edges in this special even-money case. Fractional Kelly replaces the full exposure with a fixed proportion of it. Half Kelly uses one half and quarter Kelly one quarter. Such reductions are commonly used to reduce dispersion and sensitivity to model error [5-7].

For small even-money edges, let  $e$  denote the difference between the casino win and loss probabilities and let  $\alpha$  denote the selected fraction of full Kelly.

Relative to full Kelly, the approximation implies that half Kelly retains about 75% of maximum small-edge growth, whereas quarter Kelly retains about 43.75%. The reduction in growth is less than proportional near the optimum because the objective is locally flat at its maximum. Exposure above full Kelly raises volatility and reduces expected logarithmic growth.

### 2.3. Monte Carlo design

Finite-horizon outcomes were examined with 50,000 independent terminal-bankroll simulations for each wheel and strategy over 1,000 spins. Every path began with a bankroll of one. Quarter, half, and full Kelly used the exact wheel-specific full-Kelly fraction. Separate reproducible random streams, derived from the fixed root seed 1729, were used for the terminal-distribution experiment and the spin-by-spin path figure. The fixed probabilities were held constant, so the experiment measures path variability rather than parameter uncertainty.

The recorded outputs are the median terminal bankroll, the 5th and 95th percentiles, and the probability of finishing below the initial bankroll. These statistics are more informative here than the arithmetic mean, which can be dominated by a small number of exceptionally large winning paths. With 50,000 paths, the Monte Carlo standard error for the reported loss probabilities is below approximately 0.23 percentage points.

## 3. Results and discussion

### 3.1. Analytical roulette results

Table 1 reports the exact roulette inputs and their analytical full-Kelly outputs. The American wheel supports a 5.26% exposure, compared with 2.70% for the European wheel. Although the house edge and full-Kelly fraction are numerically identical for these even-money wagers, the maximum expected log-growth rate is much smaller because only a fraction of the bankroll is exposed on each spin.

Table 1. Exact roulette inputs and full-Kelly outputs from the casino perspective

| Wheel    | Casino-win pockets | Player-win pockets | p      | House edge | Full-Kelly f | Maximum G per spin |
|----------|--------------------|--------------------|--------|------------|--------------|--------------------|
| American | 20                 | 18                 | 52.63% | 5.26%      | 5.26%        | 0.1386%            |
| European | 19                 | 18                 | 51.35% | 2.70%      | 2.70%        | 0.0365%            |

Figure 1 plots expected log growth against casino exposure for both wheels. Each curve initially rises, reaches its maximum at the wheel-specific full-Kelly fraction, and then falls as losses become increasingly damaging. Using the American full-Kelly fraction on the European game would already exceed the European optimum. Exposure must therefore be calibrated to the size of the edge, not merely to the existence of a positive edge.




Figure 1. Expected log growth as a function of casino exposure for American and European roulette (author calculations) (picture credit: original)

### 3.2. Finite-horizon simulation results

Table 2 presents the terminal-bankroll distributions after 1,000 simulated spins. For American roulette, full Kelly has the largest median and upper tail, but its 5th percentile is only 0.25 and 20.78% of paths finish below the initial bankroll. Half Kelly lowers the median from 3.87 to 2.78, yet increases the 5th percentile to 0.75 and reduces the probability of finishing below one to 10.43%. Quarter Kelly produces lower typical growth but the tightest downside distribution.

The European results have lower growth and a higher probability of underperformance because the edge is smaller. Its full-Kelly median is 1.48, while 32.42% of outcomes remain below the initial bankroll after 1,000 spins. This does not contradict positive expected logarithmic growth. The advantage compounds slowly relative to path noise, so a positive asymptotic rate does not guarantee a gain over a particular finite sample.

Figure 2 displays representative American roulette paths together with distribution bands. Full Kelly raises the median path more quickly, but its dispersion expands substantially and some paths experience deep drawdowns. Half Kelly narrows the range and raises the lower-percentile path while sacrificing part of the upper tail. This agrees with recent distributional approaches that distinguish median growth from lower-quantile or ruin-sensitive objectives [9, 10].

Table 2. Terminal bankroll after 1,000 simulated spins with an initial bankroll of one

| Wheel    | Strategy | f     | Median | 5th percentile | 95th percentile | P(final < 1) |
|----------|----------|-------|--------|----------------|-----------------|--------------|
| American | Quarter  | 1.32% | 1.82   | 0.92           | 3.60            | 7.59%        |
| American | Half     | 2.63% | 2.78   | 0.75           | 10.93           | 10.43%       |
| American | Full     | 5.26% | 3.87   | 0.25           | 59.85           | 20.78%       |
| European | Quarter  | 0.68% | 1.17   | 0.83           | 1.68            | 22.57%       |
| European | Half     | 1.35% | 1.30   | 0.64           | 2.69            | 26.31%       |
| European | Full     | 2.70% | 1.48   | 0.36           | 6.03            | 32.42%       |




Figure 2. American roulette bankroll paths under full and half Kelly; thin lines show 32 paths and thick lines show the median and 10th/90th percentiles across 3,000 paths (author simulation, root seed 1729) (picture credit: original)

### 3.3. Broader applications and limitations

Roulette is the clean benchmark because its probabilities and payouts are fixed. The same growth objective becomes more demanding when the advantage changes with the state of a game, is estimated from data, or is distributed across correlated assets. Table 3 summarises how the framework extends to blackjack, sports odds, and finance.

Table 3. Research designs for extending Kelly analysis beyond roulette

| Application | Kelly input                                                           | Proposed empirical design                                                          | Main difficulty                                      |
|-------------|-----------------------------------------------------------------------|------------------------------------------------------------------------------------|------------------------------------------------------|
| Blackjack   | Probabilities and payouts vary with deck state and rules              | Simulate shoes; compare full and fractional exposure; impose chip and table limits | State estimation, rule variation, and discrete bets  |
| Sports odds | Net odds from quoted prices; win probability from a forecasting model | Remove bookmaker margin; backtest calibrated forecasts out of sample               | Model error, market change, and weak edges           |
| Finance     | Vector returns, correlations, and portfolio constraints               | Use rolling-window log-optimal allocation with costs and drawdown analysis         | Non-stationarity, leverage, liquidity, and tail risk |

In blackjack, probabilities vary with deck composition, rules, and available actions. A card-counting system can be interpreted as a state estimator that updates the casino or player edge; practical analysis must then include discrete chip sizes, table limits, and states with no positive edge [4]. In sports betting, quoted odds determine the payoff but the win probability must be forecast. Because a small probability error can reverse a weak edge, calibration and conservative exposure are central. The treatment of multiple or overlapping wagers also requires a more general allocation model than the binary roulette calculation [7].

In finance, the single binary wager becomes a portfolio of correlated returns. The log-growth objective must be optimised subject to leverage, short-selling, liquidity, and position constraints. Recent work connects the growth-optimal direction with efficient risk-adjusted portfolios [8], while reinforcement-learning and robust distributional methods address parameter uncertainty and non-

stationarity [6, 11]. These extensions support the same practical message as the roulette results: exposure should fall when the inputs are uncertain or downside constraints are binding.

The model is not a literal operating rule for a physical casino. Real casinos accept many simultaneous bets across colours, numbers, games, and tables, and opposing customer positions can partly offset one another. Relevant exposure is therefore the net liability of a portfolio of wagers. Reserve requirements, table limits, taxes, labour costs, demand, credit risk, fraud, and regulation also affect commercial viability. In addition, the analysis assumes independent spins, fixed probabilities, frictionless reinvestment, and perfect knowledge of payoffs. The simulation isolates finite-horizon path risk; it does not model uncertain probabilities or casino-wide liabilities.

#### 4. Conclusion

The casino-side analysis shows that a structural roulette advantage supports a positive but limited growth-optimal exposure. American roulette permits a larger exposure and higher expected logarithmic growth than European roulette because its house edge is larger. Full Kelly maximises long-run compound growth under the stated model, yet the simulations show that this optimum can still generate substantial finite-horizon downside. Half and quarter Kelly surrender part of the median and upper-tail growth in exchange for improved lower-tail outcomes and reduced dispersion. A real casino may therefore rationally operate below the theoretical full-Kelly level when temporary losses, capital requirements, or model uncertainty matter. The analysis does not establish that every casino is commercially profitable; it establishes the narrower mathematical result that exposure should be scaled to the strength and reliability of the edge. Future work could replace the fixed single-wager setting with a portfolio of simultaneous liabilities and introduce uncertain, time-varying probabilities.

#### References

- [1] Ethier, S. N. (2010). *The doctrine of chances: Probabilistic aspects of gambling*. Springer. <https://doi.org/10.1007/978-3-540-78783-9>
- [2] Kelly, J. L., Jr. (1956). A new interpretation of information rate. *Bell System Technical Journal*, 35(4), 917-926. <https://doi.org/10.1002/j.1538-7305.1956.tb03809.x>
- [3] Breiman, L. (1961). Optimal gambling systems for favorable games. In J. Neyman (Ed.), *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability (Vol. 1)* (pp. 65-78). University of California Press.
- [4] Thorp, E. O. (1969). Optimal gambling systems for favorable games. *Review of the International Statistical Institute*, 37(3), 273-293. <https://doi.org/10.2307/1402118>
- [5] MacLean, L. C., Thorp, E. O., & Ziemba, W. T. (Eds.). (2011). *The Kelly capital growth investment criterion: Theory and practice*. World Scientific.
- [6] Jiang, R., Saunders, D., & Weng, C. (2022). The reinforcement learning Kelly strategy. *Quantitative Finance*, 22(8), 1445-1464. <https://doi.org/10.1080/14697688.2022.2049356>
- [7] Jacot, B. P., & Mochkovitch, P. V. (2023). Kelly criterion and fractional Kelly strategy for non-mutually exclusive bets. *Journal of Quantitative Analysis in Sports*, 19(1), 37-42. <https://doi.org/10.1515/jqas-2020-0122>
- [8] Bermin, H.-P., & Holm, M. (2024). The geometry of risk adjustments. *Decisions in Economics and Finance*, 47, 83-120. <https://doi.org/10.1007/s10203-023-00421-1>
- [9] Rubino, N., Rebholz, M., & Volpert, K. (2024). A dynamic Kelly criterion for games with a maximum payout. *The American Mathematical Monthly*, 131(7), 566-580. <https://doi.org/10.1080/00029890.2024.2351051>
- [10] Velegol, D., & Bishop, K. J. M. (2024). Analyzing sequential betting with a Kelly-inspired convective-diffusion equation. *Entropy*, 26(7), 600. <https://doi.org/10.3390/e26070600>
- [11] Sun, Y., & Zou, Z. (2025). Data-driven distributionally robust Kelly portfolio optimization based on coherent Wasserstein metrics. *JUSTC*, 55(8), 0805. <https://doi.org/10.52396/JUSTC-2024-0030>