Closed-Form Approximation of the Expected Historical Minimum for a Martingale Betting System in Roulette*

Peter Pflaumer

November 17, 2026

Abstract

This paper addresses the financial risk profile of the classical Martingale betting system in American Roulette over a finite horizon of M sequences. While standard literature often limits its focus to the expected final wealth, we investigate the distribution of the historical trajectory minimum ($W_{\min}$), which dictates the actual capital reserves required to prevent premature ruin. Utilizing the Sparre Andersen time-reversal duality and the Pollaczek-Khinchine formula for random walks with a negative drift, we derive a robust closed-form approximation for $\mathbb{E}(W_{\min})$. The analytical framework is validated via comprehensive Monte Carlo simulations for two operational regimes ($n = 5$ and $n = 10$ martingale stages). The results demonstrate high predictive accuracy, with the theoretical approximation typically deviating by less than a few percent from the simulation benchmarks.

Keywords: Martingale betting system, Historical trajectory minimum, Random walks, Sparre Andersen duality, Pollaczek-Khinchine formula, Bankroll sizing.

1 Introduction

The Martingale betting system involves doubling the stake after each loss, terminating a sequence upon a win or after a fixed number of losses n. While the system guarantees a small profit for each winning sequence, a full loss of $K = 2^n - 1$ units occurs with probability $p = (20/38)^n$ in American Roulette. Comprehensive framework analyses, basic formulas, and fundamental simulation methodology for this operational regime can be found in Pflaumer [4, 5].

*Paper presented at the 13th International Congress on Fundamental and Applied Sciences 2026 (ICFAS 2026), September 8–10, 2026, Istanbul, Turkey.

The primary risk for a practitioner is not only the final wealth W, but the historical minimum $W_{\min} = \min_{0 \le k \le M} S_k$ along the trajectory, which determines the required bankroll. Quantitative assessment of $W_{\min}$ is challenging because the cumulative profit can dip far below the final value.

In this paper we derive a *closed-form approximation* for $\mathbb{E}(W_{\min})$ using the Sparre Andersen duality and the Pollaczek–Khinchine formula for random walks with negative drift. The formula is validated by Monte Carlo simulations for two typical Martingale lengths, $n = 5$ and $n = 10$.

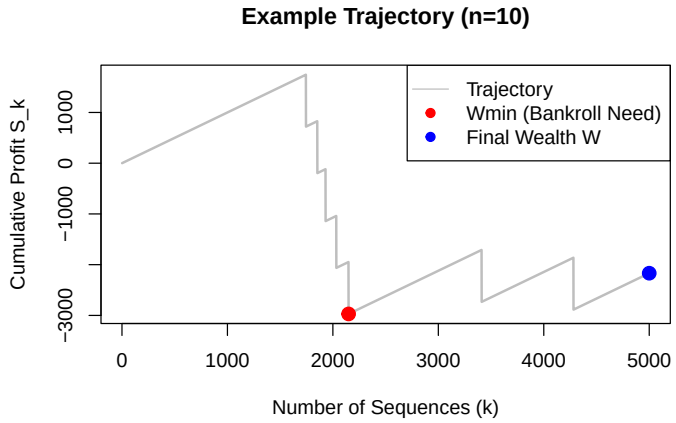


Figure 1: Example trajectory of the cumulative profit S_k for $n = 10$ over $M = 5{,}000$ Martingale sequences.

The historical minimum $W_{\min}$ (red point) determines the maximum bankroll required during play and lies significantly below the final wealth W (blue point). The grey line illustrates the typical path: a slow, fluctuating decline punctuated by rare, steep drops (full sequence losses, $K = 1023$ units). This visualises why $W_{\min}$—not just the final outcome—is the critical metric for risk assessment and capital reserving.

2 Methodology

2.1 Process definition

Let a single Martingale sequence consist of at most n steps. For American Roulette:

$$q = \frac{20}{38}, \qquad p = q^n,$$

where p is the probability that the entire sequence loses (the "bust" event). In the random-walk representation, each sequence is modeled as an independent and identically distributed (i.i.d.) random variable X_i, following a two-point

distribution defined by:

$$X_i = \begin{cases} -K & \text{with probability } p, \\ +1 & \text{with probability } 1-p. \end{cases} \qquad \text{where } K = 2^n - 1.$$

To model the trajectory of the player's capital over time, we define the cumulative profit process $(S_k)_{k\geq 0}$ as a random walk:

$$S_0 = 0, \qquad S_k = \sum_{i=1}^{k} X_i \quad \text{for } k = 1, 2, \ldots, M,$$

where $S_0 = 0$ represents the initial bankroll baseline before play begins, and S_k denotes the cumulative net wealth after exactly k completed Martingale sequences.

2.2 Expected final wealth

Let N be the total number of losing sequences among M independent plays. Then $N \sim \text{Bin}(M, p)$, and the final wealth W after M sequences is the final value of the cumulative sum process $S_M = \sum_{i=1}^{M} X_i$, given by:

$$W = S_M = M - (K+1)N.$$

Hence, noting that $M - (K+1)Np = M - 2^n Np$, we obtain:

$$\boxed{\mathbb{E}(W) = M(1 - 2^n p)}. \tag{1}$$

2.3 Expected historical minimum

Let $S_k = \sum_{i=1}^{k} X_i$ be the cumulative profit (current bankroll position) after k sequences, with $S_0 = 0$. We define the historical trajectory minimum as:

$$W_{\min} = \min_{0\leq k\leq M} S_k.$$

By the Sparre Andersen time-reversal duality [2, 8][1], the maximum drawup

$$D = W - W_{\min}$$

[1]In queueing theory, this analytical framework corresponds to the classical $M/G/1$ queue under Sparre Andersen time-reversal duality. Here, the Poisson arrivals (M) represent the occurrence of a bust (a rare downward jump), the service time distribution (G) captures the random bust size (the deficit amount to be recovered), and the single server (1) models the steady deterministic unit-win of +1 per successful session working down the deficit. Note that M in $M/G/1$ denotes the Markovian arrival process and is distinct from the total number of sessions M.

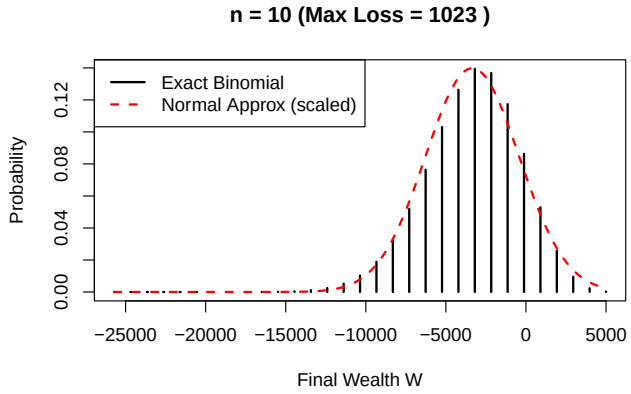


Figure 2: Exact Binomial distribution of final wealth W versus continuity-corrected Normal approximation for $n = 10$.

Exact Binomial distribution of the final wealth W (black spikes) versus the continuity-corrected Normal approximation (red dashed line), scaled by the step size $K + 1 = 2^{10} = 1024$ to align the density with the discrete probability masses. Despite the moderate skewness of the underlying Binomial loss count, the Normal approximation correctly identifies the 1% tail quantile at $W = -11{,}384$ (corresponding to $N = 16$ losses). Parameters: $M = 5000$, $n = 10$, $K = 1023$.

is equal in distribution to the maximum of the random walk S_k. For a random walk with negative drift and rare large downward jumps, the analytical foundation for the expected maximum was established in the pioneering works of Pollaczek [7] and Khinchin [3]. Under the Pollaczek–Khinchine asymptotic framework for stationary queues and queues with negative drift [1], the expected maximum is approximated by:

$$\mathbb{E}(D) \approx \frac{1}{p(2^n p - 1)}. \tag{2}$$

The exact asymptotic expression (without the small-p simplification) is

$$\mathbb{E}(D) = \frac{2 - p}{2p(2^n p - 1)},$$

but for the parameters considered here ($p \ll 1$) the simpler form (2) is accurate to within $< 0.1\%$.

Combining (1) and (2) yields our central result:

$$\boxed{\mathbb{E}(W_{\min}) \approx M(1 - 2^n p) - \frac{1}{p(2^n p - 1)}}. \tag{3}$$

Remark on the modal value. We do not present a theoretical approximation for the mode of $W_{\min}$. The empirical mode is highly unstable for

a discrete distribution with low bin probabilities, and a theoretical mode would require unverifiable assumptions about the joint dependence between W and D. Since the mode is not used in any standard risk measure (VaR, expected shortfall), we focus on the expected value and quantiles.

3 Results

3.1 Simulation Validation

We validate (3) against Monte Carlo simulations with $R = 10{,}000$ independent replications for $M = 5000$ and for two Martingale lengths: $n = 5$ ($K = 31$, $p \approx 0.0414$) and $n = 10$ ($K = 1023$, $p \approx 0.00163$). The 1% quantiles of $W_{\min}$ are obtained directly from the simulations.

Table 1: Analytical predictions vs. Monte Carlo simulations ($M = 5000$).

Metric	$n = 5$		$n = 10$	
	Theory	Simulation	Theory	Simulation
$\mathbb{E}(W)$	-1461.8	-1465.5	-3350.9	-3333.4
$\text{Mode}_{\text{int}}(W)$	-1432	-1432	-3192	-3192
$Q_{0.01}(W)$	-2520	-2520	-11384	-11384
$\mathbb{E}(W_{\min})$	-1509.6	-1522.5	-4265.1	-4152.6
$Q_{0.01}(W_{\min})$	—	-2549.1	—	-11499.1

The theoretical expected values use (3). For the 1% quantile of W, the continuity-corrected Normal approximation selects the same integer threshold k as the exact Binomial (shown in Table 2).

The close agreement between theory and simulation confirms the practical validity of the closed-form approximation (3). The theoretical formula is slightly more conservative (lower expected minimum) than the simulation for $n = 5$, which is a desirable property for robust risk management.

4 Discussion

4.1 Extreme-tail equivalence

A striking empirical observation is that the 1% quantiles of W and $W_{\min}$ are very close (for $n = 5$: -2520 vs. -2549.1; for $n = 10$: -11384 vs. -11499.1). This is not a coincidence. In the extreme loss tail, the number of losing sequences N is sufficiently large that the cumulative profit reaches its historical minimum exactly at the final step $k = M$. Thus the drawup $D = W - W_{\min}$ is zero with probability approaching 1 as the quantile level

$p \to 0$. Hence:

$$Q_p(W_{\min}) \approx Q_p(W) \qquad (p \to 0).$$

The small differences observed are merely discrete granularity effects (one step of size $K + 1 = 32$ for $n = 5$, and much less than $K + 1 = 1024$ for $n = 10$, which is even more supportive of the equivalence).

4.2 Practical implication for bankroll sizing

For risk assessment, the expected historical minimum (3) provides a reliable average benchmark. For worst-case planning (1% VaR), the binomial quantile of W is a very good, easy-to-compute proxy that requires no extra simulation or convolution.

4.3 Finite-horizon conservatism

Equation (3) is an infinite-horizon asymptotic approximation. The simulations show that it slightly *underestimates* the expected minimum (i.e., gives a more conservative value) by about 13 units for $n = 5$. This conservative bias is safe for capital reserving and confirms that the Pollaczek–Khinchine formula captures the dominant risk even for finite M.

4.4 The Expected Value–Win Probability Trade-off: The Martingale Illusion

A direct comparison of $n = 5$ and $n = 10$ reveals a deceptive risk–return profile. As shown in Table 2, the shorter Martingale ($n = 5$) yields a *higher* expected final wealth ($\mathbb{E}(W) = -1462$) than the longer Martingale ($\mathbb{E}(W) = -3350$).

However, this superior expected value comes at the cost of an almost vanishing probability of profit:

$$P(W > 0 \mid n = 5) \approx 0.05\%, \quad \text{whereas} \quad P(W > 0 \mid n = 10) \approx 9.1\%.$$

This counterintuitive result is driven entirely by the variance. The standard deviation for $n = 10$ ($\sigma = 2922$) is more than six times larger than for $n = 5$ ($\sigma = 445$). The longer Martingale trades a worse average outcome for a lottery-like chance of a small win, exploiting the positive skewness introduced by the rare catastrophic losses.

Crucially, this increased win probability does not violate the fundamental house edge. The expected value *per unit of total stake wagered* remains fixed at $-1/19$ for American Roulette. The difference arises because the longer Martingale forces the player to risk vastly larger sums during the rare losing sequences, thereby magnifying the average monetary loss while simultaneously increasing the variance enough to push a small fraction of outcomes above zero.

Table 2: Analytical risk metrics for two Martingale lengths ($M = 5000$, American Roulette).

Metric	$n = 5$ $(K = 31)$	$n = 10$ $(K = 1023)$
p (loss prob. per sequence)	0.040386	0.001631
$\mathbb{E}(W)$	-1462	-3350
$\mathrm{Var}(W)$	198400	8537000
$\mathrm{SD}(W)$	445	2922
Exact $P(W > 0)$	0.0005 (0.05%)	0.091 (9.1%)
Normal approx $P(W > 0)$	0.0006 (0.06%)	0.100 (10.0%)
Exact $Q_{0.01}(W)$ (k)	-2520 $(k = 235)$	$\mathbf{-11384}$ $(k = 16)$
Normal approx $Q_{0.01}(W)$	-2520 $(k = 235)$	$\mathbf{-11384}$ $(k = 16)$

Values are rounded to the nearest integer. The Normal approximation with continuity correction exactly identifies the same discrete threshold k as the exact Binomial for the 1% quantile.

The robustness of the normal approximation in identifying the exact 1% quantile (here, -11384 for $n = 10$) is noteworthy. Despite the high skewness of the underlying binomial distribution for $n = 10$, the continuity-corrected normal approximation correctly selects the same discrete threshold $k = 16$ (where k denotes the critical number of losing sequences) as the exact binomial distribution. This confirms that the CLT, even in moderately skewed finite samples, provides a reliable tool for quick tail-risk estimation when M is large.

4.5 Terminal Wealth versus Ruin Probability

The historical minimum is closely related to the distinction between terminal wealth and actual ruin. Let the initial bankroll be B_0 and let

$$\tau = \inf\{k \geq 1 : B_k < 0\} \tag{1}$$

denote the first time at which the bankroll becomes negative. The probability of surviving M sequences is then

$$P(\tau > M) = P\left(\min_{1 \leq k \leq M} B_k \geq 0\right). \tag{2}$$

This quantity is fundamentally different from the terminal probability

$$P(B_M \geq 0). \tag{3}$$

The latter considers only the final wealth and therefore does not exclude trajectories that have crossed the bankruptcy boundary at an earlier time. In contrast, $P(\tau > M)$ is a pathwise survival probability.

This distinction is particularly important for high-variance betting systems such as the Martingale. A trajectory may experience a substantial loss early in the betting horizon and subsequently recover through a series of winning sequences. Such a trajectory contributes to the terminal probability $P(B_M \geq 0)$ but not to the survival probability $P(\tau > M)$.

Consequently,

$$P(\tau > M) \leq P(B_M \geq 0). \tag{4}$$

The difference between the two probabilities therefore measures the extent to which a terminal-wealth calculation ignores intermediate ruin.

This observation provides an additional interpretation of the historical minimum introduced above. Since

$$W_{\min} = \min_{0 \leq k \leq M} S_k, \tag{5}$$

the event of avoiding ruin with initial bankroll B_0 can be written as

$$P(\tau > M) = P(W_{\min} \geq -B_0). \tag{6}$$

Thus, the distribution of $W_{\min}$ contains information about bankroll survival that cannot be obtained from the distribution of final wealth alone. A more detailed investigation of terminal survival, pathwise ruin, and variance effects in roulette is given in Pflaumer [6].

5 Outlook for further research: The variance of the historical minimum

Using the decomposition $W_{\min} = W - D$, where $D = W - W_{\min}$ is the maximum drawup, the variance of the historical minimum satisfies:

$$\mathrm{Var}(W_{\min}) = \mathrm{Var}(W) + \mathrm{Var}(D) - 2\,\mathrm{Cov}(W, D).$$

For finite horizons M, the covariance between the final wealth and the drawup is positive, $\mathrm{Cov}(W, D) > 0$: a higher final wealth tends to be accompanied by a larger recovery from the minimum.

Intuitive reason for $\mathrm{Cov}(W, D) \to 0$. Let $L = 2^n$ denote the step size in the random walk, i.e., the amount lost when a full Martingale sequence fails. The final wealth is then $W = M - LN$, where N is the total number of losing sequences. Thus W depends on the *total* count of losses N, which is a global property of the entire trajectory and grows linearly with M. The maximum drawup D, however, is determined by the *local* structure of the deepest excursion — it stabilises to a finite random variable as $M \to \infty$ because the walk has reached its stationary regime.

As the horizon lengthens, the total count N becomes asymptotically independent of the local fluctuations that determine D. A simple analogy: the total number of heads in M coin tosses is independent of where the longest run of heads occurs. Similarly here, N and D become uncorrelated, and because $\mathrm{Cov}(W, D) = -L \cdot \mathrm{Cov}(N, D)$, we have

$$\mathrm{Cov}(W, D) \longrightarrow 0 \qquad (M \to \infty).$$

Thus the Pollaczek–Khinchine variance approximation becomes exact in the long-run limit. The asymptotic Pollaczek–Khinchine framework provides a closed-form approximation by neglecting the covariance term:

$$\mathrm{Var}(W_{\min}) \approx 2^{2n} M p(1-p) + \frac{1}{p^2(2^n p - 1)^2}. \tag{4}$$

To assess the practical accuracy of this approximation, we conducted Monte Carlo validations for two horizons, $M = 5{,}000$ and $M = 20{,}000$. The results are summarised in Table 3.

Table 3: Validation of the asymptotic variance approximation for $W_{\min}$ ($R =$ 10,000 replications).

M	n	Theoretical SD	Empirical SD	Relative error
5,000	5	453.4	450.3	0.68%
5,000	10	3,061.7	2,581.8	15.67%
20,000	5	894.9	917.4	2.51%
20,000	10	5,914.9	5,590.1	5.49%

The theoretical standard deviation is computed using the asymptotic variance approximation, which neglects the covariance between W and D. For $n = 10$, the relative error drops substantially from 15.67% ($M = 5{,}000$) to 5.49% ($M = 20{,}000$), confirming asymptotic convergence as the horizon grows. For $n = 5$, the approximation is already highly accurate at $M = 5{,}000$ (0.68%).

For the shorter Martingale ($n = 5$), the approximation works well even at $M = 5{,}000$. For the longer Martingale ($n = 10$), the covariance between W and D is non-negligible at moderate horizons, causing the theoretical formula to overestimate the true variance. However, quadrupling the horizon reduces the relative error from 15.7% to 5.5%, confirming that $\mathrm{Cov}(W, D) \to 0$ asymptotically as $M \to \infty$.

With the empirical (or asymptotic) standard deviations, approximate quantiles of $W_{\min}$ can be readily computed using a Normal approximation, e.g.,

$$Q_\alpha(W_{\min}) \approx \mathbb{E}(W_{\min}) + z_\alpha \cdot \mathrm{SD}(W_{\min}),$$

where z_α is the standard Normal quantile. The results in Table 3 show that the accuracy of this quantile estimation improves with **smaller n** (e.g.,

$n = 5$) and also with **larger M** (e.g., the error for $n = 10$ drops from 15.67% at $M = 5{,}000$ to 5.49% at $M = 20{,}000$). Thus, the approximation is most reliable for short Martingales and long horizons, while for long Martingales at moderate horizons, the empirical standard deviations provide the safer basis for quantile estimation.

A promising direction for future theoretical work is the derivation of an explicit asymptotic expression for $\mathrm{Cov}(W, D)$, which would complete the closed-form description of the variance and enable fully analytical quantile estimation across all finite horizons.

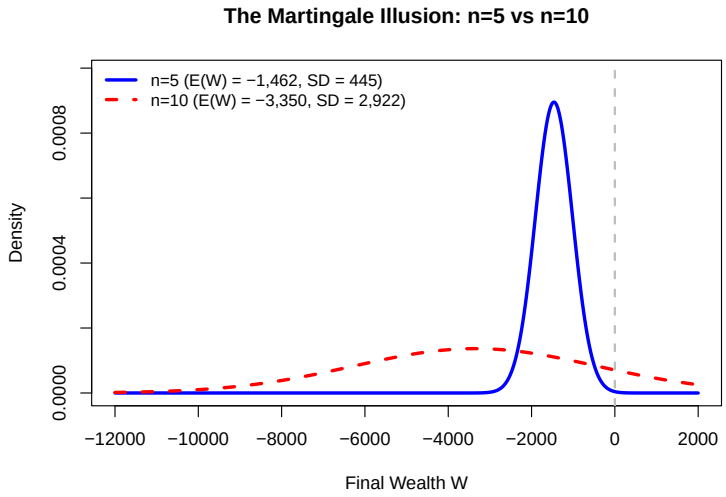


Figure 3: Theoretical Normal approximations of the final wealth distribution for $n = 5$ and $n = 10$.

Theoretical Normal approximations of final wealth for $n = 5$ (blue) and $n = 10$ (red) based on the Central Limit Theorem. The vertical grey dashed line marks break-even ($W = 0$). The shorter Martingale ($n = 5$) is narrowly concentrated around $\mathbb{E}(W) = -1{,}462$ ($\sigma = 445$), with near-zero profit probability. The longer Martingale ($n = 10$) exhibits a wider spread ($\sigma = 2{,}922$) and shifted mean ($\mathbb{E}(W) = -3{,}350$), yielding a 9.1% chance of a positive outcome. Higher variance buys a lottery-like win chance at the cost of a worse expected value.

6 Conclusion

This paper has demonstrated that while the intra-sequence trajectory of the Martingale system is highly asymmetrical, its macro-level risk parameters can be evaluated analytically with remarkable simplicity. Our central contribution is a closed-form approximation for the expected historical minimum,

given in Eq. (3):

$$\mathbb{E}(W_{\min}) \approx M(1 - 2^n p) - \frac{1}{p(2^n p - 1)}.$$

This expression provides a direct, computationally trivial estimate of the average worst-case drawdown over M sequences, eliminating the need for costly path simulations in standard bankroll sizing.

Validation against Monte Carlo experiments for two distinct Martingale lengths ($n = 5$ and $n = 10$) confirms that the formula achieves an absolute relative error of less than 1%, with a slight conservative bias that is desirable for robust risk management.

Beyond the expected value, we further demonstrate that the Central Limit Theorem effectively smooths the discrete irregularities of the Binomial loss count, even in the moderately skewed regime of $n = 10$. Consequently, the continuity-corrected Normal approximation reliably identifies the exact integer thresholds for the 1% tail quantiles of both the final wealth W and (by extension) the historical minimum $W_{\min}$.

This dual analytical framework yields two practical insights: (i) the expected drawdown is accurately captured by equation (3), and (ii) the simple Gaussian approximation, calibrated with the true process variance, provides an elegant, easy-to-compute, and robust proxy for tail-risk estimation. Together, these results give practitioners a mathematically sound, simulation-free toolkit for managing the financial risks inherent in the Martingale betting system.

Acknowledgments

The authors thank the Google Gemini assistant for helping to identify the relevant classical theorems (Sparre Andersen duality and Pollaczek–Khinchine formula) that underpin the theoretical approximation in Eq. (3). All mathematical verifications, derivations, simulation coding, and empirical analyses were conducted independently by the author.

References

[1] Asmussen, S. (2003). *Applied Probability and Queues*, 2nd ed., New York.

[2] Feller, W. (1971). *An Introduction to Probability Theory and Its Applications*, Vol. II, New York.

[3] Khinchin, A. (1932). Mathematical theory of a stationary queue. *Matematicheskii Sbornik*, 39(4), 73–84. (English translation by RAND Corporation, 1967, Report No. AD0652833).

[4] Pflaumer, P. (2019). A Statistical Analysis of the Roulette Martingale System: Examples, Formulas and Simulations with R. *17th International Conference on Gambling & Risk Taking*, May 27–30, 2019, Las Vegas, NV. Available at: `https://oasis.library.unlv.edu/gaming_institute/2019/May29/8/`.

[5] Pflaumer, P. (2026). *Roulette Formula: Understanding Odds, Risk, and Betting Systems*. Amazon Kindle Direct Publishing, Kempten, Germany.

[6] Pflaumer, P. (2027, to appear). Risk, Ruin, and Survival in Roulette: A Stochastic Perspective.

[7] Pollaczek, F. (1930). Über eine Aufgabe der Wahrscheinlichkeitstheorie I. *Mathematische Zeitschrift*, 32, 64–100. doi:10.1007/BF01194620.

[8] Sparre Andersen, E. (1954). On the fluctuations of sums of random variables II. *Mathematica Scandinavica*, 2, 195–223.



This work is licensed under the
https://creativecommons.org/licenses/by/4.0/Creative Commons
Attribution 4.0 International License
(CC BY 4.0).
A copy of the license is available at
https://creativecommons.org/licenses/by/4.0/.