Samurai Survival Statistics: a Monte-Carlo approach to game-theoretical sudden death interactions

Paulo Garcia
Carnegie Mellon-KMITL Thailand Program, CMKL University
Bangkok, Thailand
paulo@cmkl.ac.th

I have trained in the way of strategy since my youth, and at the age of thirteen I fought a duel for the first time. My opponent was called Arima Kihei, a sword adept of the Shinto ryu, and I defeated him. At the age of sixteen I defeated a powerful adept by the name of Tadashima Akiyama, who came from Tajima Province. At the age of twenty-one I went up to Kyoto and fought duels with several adepts of the sword from famous schools, but I never lost.

Miyamoto Musashi, Go Rin No Sho

***Abstract*—Sword-duels are interesting, from a game-theoretical perspective, as they correspond to sudden-death interactions between players (i.e., defeat means removal from the game), and where victory/defeat depend on not just skill, but also luck. Analyzing probabilities of victory streaks, given a certain level of self and others' skill, is thus relevant for any application domain that can be modeled in the same way.**

This paper takes inspiration from Miyamoto Musashi, famously undefeated for 61 duels, and implements a Markov-chain Monte-Carlo simulation approach to evaluate this scenario. Results suggest that a 61 victory streak can be probabilistically observed when skill level is roughly 6.5 times that of the average duelist, and with 95% confidence when skill level is roughly 1000 times that of the average duelist.

More generally, this paper provides a method for determining chances of victory streaks in game-theoretical sudden-death encounters, when both "skill" and "luck" contribute to the outcome of the encounter. Specific scenarios can be modeled by modifying the utilized Markov chain and adjusting sampled distributions as required.

I. Introduction

Miyamoto Musashi (1584 - 1645) [1] was a ronin that was famously undefeated for 61 consecutive duels. Sword duels are particularly interesting, from a game-theoretical perspective, as they represent a combination of skill and luck (i.e., randomness) [2]: a more skilled fighter is more likely to, but not certain to, win a duel against a less skilled opponent. In the case of duels to the death, each duelist (player) can achieve a string of victories, greater than or equal to 0, before being removed from the game.

This paper examines this scenario. Specifically, we show how to mathematically model both skill and luck, and examine how this game evolves computationally, using Monte-Carlo methods [3]. Our set up allows us to examine two related research questions:

Fig. 1. Musashi and Kojiro in battle (statue: Island of Ganryujima) [8]

1) Given an upper bound on skill, what is the expected number of victories the highest skilled player can expect, given that skill follows a normal distribution?
2) If skill were boundless, how skilled would the most skilled player have to be, to achieve a given number of victories (we examine 61, in honor of Musashi) with 95% confidence?

This scenario is relevant not just to sword duels, but generally for any game with discrete competitions between two players, where defeat means removal from the game. "Skill" is equivalent to player properties in other settings (e.g., money, fame, athleticism, expertise, medical precautions) and "luck" is equivalent to external variables (e.g., weather conditions, stock market state). Example application scenarios include financial endeavors where "defeat" (e.g., failing to purchase a commodity, obtain a bid) means irreparable loss of capital [4]; social interactions where "defeat" (e.g., failing to woo a potential mate, perform some rite of passage) means being shunned from the social circle [5]; government covert intelligence where "defeat" (e.g., failing to obtain sensitive intelligence that falls into enemy hands) means irreparable loss of life and/or sovereignty [6]; or pandemic prevention where "defeat" (e.g., being infected by an airborne disease such as COVID-19) may prove lethal [7]. Throughout the remainder of this paper, we continue to use the sword duel analogy and vocabulary.

The remainder of this paper is organized as follows: Section II describes our mathematical model. Section III describes the implementation of the Monte-Carlo simulation. Section IV presents our computational results, with Section V concluding this paper. All code used in the experiments is open-source, available here[1].

II. The Mathematical Model

A. Sword duels

Sword duels are discrete events. They happen sequentially, and only between two duelists (we do not model 3 or more duelist engaged in all-vs-all, nor N-vs-M engagements). A duelist i has an associated skill S_i. If a duelist i engages duelist j in a duel, we model $P(I)$, the probability that i wins the duel, as:

$$P(I) = \frac{S_i}{S_i + S_j} \tag{1}$$

Notice that $P(I) + P(J) = 1$, i.e., we assume draws are not possible. This can be interpreted as: "a duelist of skill $S_i = k \times S_j, k \in \mathbb{N}$ will defeat duelists of skill S_j, k times out of $k+1$ duels". Our model generates a random number, sourced from a uniform distribution, to decide victory in duels, thus modeling luck.

This probabilistic victory model is an assumption that can of course be changed depending on the scenario; e.g., skill can increase chances of victory in a quadratic or exponential way, rather than in a linear way, but this model can be used without loss of generality.

B. Duelists

Depending on which research question we are interested in answering, we model duelists' skill as either a two-sided truncated normal distribution (skill has both a lower and upper bound) or as a left-truncated (lower-tail) normal distribution (skill has a lower bound, but no upper bound). I.e., we are assuming that there is a lower bound on skill for someone to engage in sword duels, denoted by S_0, and potentially a maximum skill level, denoted by S_{max}.

The probability density function (PDF) for skill S with mean μ and variance σ^2, left-truncated at S_0 and right-truncated at S_{max}, is given by:

$$f(S; \mu, \sigma, S_0, S_{max}) = \frac{1}{\sigma} \frac{\phi(\frac{S-\mu}{\sigma})}{\Phi(\frac{S_{max}-\mu}{\sigma}) - \Phi(\frac{S_0-\mu}{\sigma})} \tag{2}$$

for $S_0 < S < S_{max}$, and 0 otherwise, where:

$$\phi(\xi) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}\xi^2\right) \tag{3}$$

is the probability density function of the standard normal distribution and

$$\Phi(x) = \frac{1}{2}\left(1 + \mathrm{erf}(x/\sqrt{2})\right) \tag{4}$$

[1]https://github.com/paulofrgarcia-cmkl/samurai_survival_statistics

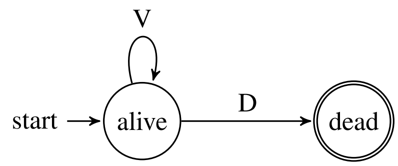

Fig. 2. The sword duelist's state space as a Markov chain. Each duel updates the player's state according to victory ("V") or defeat ("D"). V and D are generated through a Monte-Carlo approach.

is its cumulative distribution function [9].

In the case of no upper bound on skill, the probability density function for skill S is given by:

$$f(S; \mu, \sigma, S_0) = \frac{1}{\sigma} \frac{\phi(\frac{S-\mu}{\sigma})}{1 - \Phi(\frac{S_0-\mu}{\sigma})} \tag{5}$$

for $S_0 < S$, and 0 otherwise, since

$$\Phi\left(\frac{\infty-\mu}{\sigma}\right) = 1, \tag{6}$$

III. Monte-Carlo simulation methodology

Our hypothetical Musashi is modeled as an extremely simple Markov chain [10], depicted in Fig. 2. The duelist begins in the "alive" state. Every duel updates its state, in function of victory or defeat, generated as per the Monte-Carlo methods described below. More complex state spaces and update functions can be implemented using the same methodology.

A. Upper/lower-bounded skill

This method answers the research question "Given an upper bound on skill, what is the expected number of victories the highest skilled player can expect, given that skill follows a normal distribution?".

To evaluate this question, we set up a Monte-Carlo simulation where Musashi skill S_i is heuristically fixed at 100. A two-sided truncated normal distribution $f(S; \mu, \sigma, S_0, S_{max})$ is sampled to generate opponent skill level. We set minimal skill as 1, and mean as $\frac{100-1}{2}$, such that $f(S; 49.5, \sigma, 1, 100)$. Once an opponent skill level S_j has been generated, a duel is performed, sampling a uniform distribution between 0 and $S_i + S_j$. If the result is less than or equal to S_j, defeat is determined; else, victory is determined (this corresponds to the probability in Equation 1).

For every evaluated value of σ^2 (i.e, examining how the variance affects the result), we perform 1000 iterations. From those, we report the average number of victories, the minimum number of victories that occurred 50% of the iterations, the minimum number of victories that occurred 95% of the iterations, and the highest observed number of victories. σ^2 is varied between 0 and μ, in increments of $0.1 \times \mu$.

B. Lower-bounded skill

If skill were boundless, how skilled would the most skilled player have to be, to achieve a given number of victories (we examine 61, in honor of Musashi) with 95% confidence?

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 22 November 2022 doi:10.20944/preprints202211.0419.v1

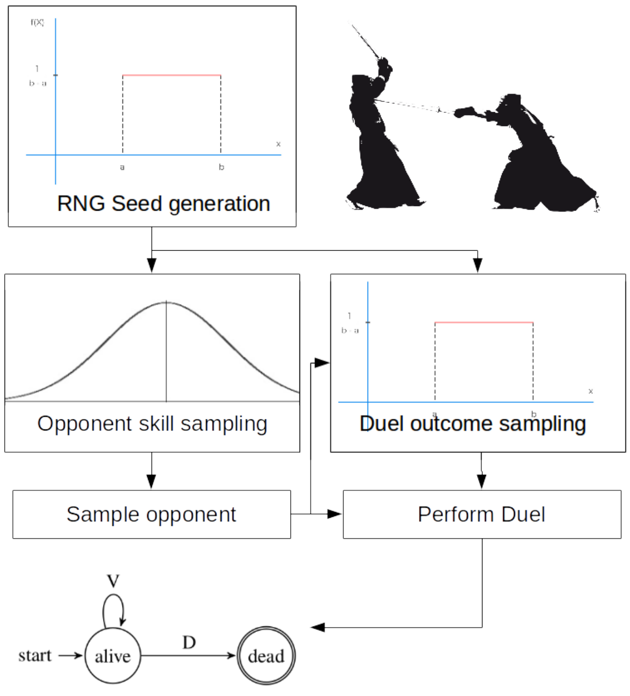

Fig. 3. Monte-Carlo simulation: depiction of 1 iteration, for given values of distribution parameters (part of image reproduced from [12])

To evaluate this question, we set up a Monte-Carlo simulation where Musashi skill is varied, with an initial value of 100. Mean is fixed at 49.5, so results can be compared with the previous experiment. As before, σ^2 is varied between 0 and μ, in increments of $0.1 \times \mu$. For each value of σ^2 and Musashi skill, 1000 iterations are performed. For a given value of σ^2, Musashi skill is increased by 10, until 61 victories are observed over 95% of the time. For each pair of σ^2, Musashi skill, we report the lowest skill level that obtained at least one chain of length 61, the lowest skill level that obtained a chain of length 61 over 95% of the time, and the average victories at that skill level.

C. Implementation

The implemented Monte-Carlo simulation is depicted in Fig. 3. Distribution sampling is performed using the GNU Scientific Library (GSL [11]). Noteworthy details: when normal distributions are used, the associated Random Number Generators (RNGs) have to be re-initialized (in our implementation, by re-seeding them) across iterations, to guarantee statistical independence across results. Were RNGs not re-initialized, results from one iteration could pollute results from subsequent iterations; e.g., if sampled numbers from a normal distribution happen to come from the lower part of the distribution in one iteration, a subsequent iteration is more likely to sample from the upper part of the distribution, without re-initialization (this is not observed when sampling from a uniform distribution). We use an extra (uniform) RNG to seed normal distributions across iterations.

IV. Results

Table I depicts results for the two-sided truncated normal distribution experiment, i.e., with an upper bound on skill.

TABLE I
VICTORY RESULTS FOR FIXED MUSASHI SKILL, ASSUMING TWO-WAY TRUNCATED NORMAL DISTRIBUTION.

σ^2	Average	guaranteed 50%	guaranteed 95%	Longest
4.95	2.017	1	0	15
9.90	1.988	1	0	19
14.85	2.047	1	0	16
19.80	2.100	1	0	16
24.75	1.962	1	0	17
29.70	1.899	1	0	15
34.65	1.901	1	0	18
39.60	1.755	1	0	20
44.55	1.787	1	0	22

For a range where the hypothetical Musashi is 100 times as good as the least skilled duelist, and twice as good as the average duelist, we observe that 1 victory can be observed with 50% confidence, but no victories can be guaranteed with 95% confidence. On average, we win slightly less than 2 duels in a row, with a highest observable streak of 22 victories.

Table II depicts results for the one-sided (left) truncated normal distribution, i.e., with a lower, but not upper, bound on skill.

TABLE II
VICTORY RESULTS FOR VARIED MUSASHI SKILL, ASSUMING ONE-WAY TRUNCATED NORMAL DISTRIBUTION. S_{61} IS THE SKILL VALUE CORRESPONDING TO AT LEAST ONE OBSERVED 61 VICTORY STREAK; $S_{95\%}$ IS THE SKILL VALUE FOR 61 VICTORIES WITH 95% CONFIDENCE.

σ^2	S_{61}	Average at $S_{95\%}$	$S_{95\%}$	Longest at $S_{95\%}$
4.95	260	975.9	48120	8997
9.90	330	988.3	50080	10454
14.85	290	950.6	51190	9618
19.80	260	1032.2	49600	10427
24.75	300	929.5	46040	8239
29.70	300	1003.5	52610	10115
34.65	400	998.9	59600	9638
39.60	360	1007.0	59070	9637
44.55	400	1046.8	60110	9542

We observe the first streak of 61 victories when the hypothetical Musashi skill is, on average, 322.2; i.e., 322 times as good as the weakest duelist, and roughly 6.5 times as good as the average duelist. To obtain the same streak with 95% confidence, the hypothetical Musashi skill must be 52935.5, i.e., a thousand times better than the average duelist.

V. Conclusions

We have evaluated victory streak probabilities for sword duel survival, using a Markov-chain Monte-Carlo simulation approach. Results suggest that a 61 victory streak (chosen in honor of Musashi's victory streak) can be probabilistically observed when skill level is roughly 6.5 times that of the average duelist, and with 95% confidence when skill level is roughly 1000 times that of the average duelist.

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 22 November 2022 doi:10.20944/preprints202211.0419.v1

More generally, this paper provides a method for determining chances of victory streaks in game-theoretical sudden-death encounters, when both ”skill” and ”luck” contribute to the outcome of the encounter. Specific scenarios can be modeled by modifying the utilized Markov chain and adjusting sampled distributions as required.

As a final remark: applying statistical probabilities to events in the past is, at best, a subversion of the field and, at worst, an epistemological challenge to the meaning of probability. Thus, we do not attempt to answer this question, but we do leave it for the reader to ponder, if they so desire: ”how good was Miyamoto Musashi?”

References

[1] K. Tokitsu, *Miyamoto Musashi: His life and writings*. Shambhala Publications, 2006.

[2] T. Nadelhoffer, “Skill, luck, control, and intentional action,” *Philosophical Psychology*, vol. 18, no. 3, pp. 341–352, 2005.

[3] D. P. Kroese, T. Brereton, T. Taimre, and Z. I. Botev, “Why the monte carlo method is so important today,” *Wiley Interdisciplinary Reviews: Computational Statistics*, vol. 6, no. 6, pp. 386–392, 2014.

[4] E. F. Fama and K. R. French, “Luck versus skill in the cross-section of mutual fund returns,” *The journal of finance*, vol. 65, no. 5, pp. 1915–1947, 2010.

[5] S. Krull, M. Pelster, and P. Steinorth, “Skill, effort, luck: The impact of rankings on risk-taking in a social setting,” *Available at SSRN 4223297*, 2022.

[6] J. B. Bell, “The armed struggle and underground intelligence: an overview,” *Studies in Conflict & Terrorism*, vol. 17, no. 2, pp. 115–150, 1994.

[7] B. Flyvbjerg, “The law of regression to the tail: How to survive covid-19, the climate crisis, and other disasters,” *Environmental Science & Policy*, vol. 114, pp. 614–618, 2020.

[8] RoninjapanWiki, “Island of ganryujima,statue of musashi and kojiro battle.” https://www.maxpixel.net/Tsuki-Samurai-Shock-Japan-Engraving-Fight-Kendo-5381023.

[9] M. T. Khasawneh, S. R. Bowling, S. Kaewkuekool, and B. R. Cho, “Tables of a truncated standard normal distribution: A singly truncated case,” *Quality Engineering*, vol. 17, no. 1, pp. 33–50, 2004.

[10] W.-K. Ching and M. K. Ng, “Markov chains,” *Models, algorithms and applications*, 2006.

[11] M. Galassi, J. Davies, J. Theiler, B. Gough, G. Jungman, P. Alken, M. Booth, F. Rossi, and R. Ulerich, *GNU scientific library*. Network Theory Limited Godalming, 2002.

[12] M. Pixel, “Tsuki samurai shock japan engraving fight kendo,” https://en.wikipedia.org/wiki/Miyamoto_Musashi.