Forecasting the Next Speedrun World Record for a Video Game

A comparison of exponential and Weibull models
for the waiting time between records

Mika Han

Contents

1 Introduction and Rationale **3**

2 Aim and Approach **3**

3 Exploration **5**

3.1 Definitions and Data Collection 5

3.2 Modelling Speedrun World Record Progression 5

3.2.1 Exponential Regression Model 7

3.2.2 Logarithmic Linearisation 8

3.2.3 Least-Squares Estimation 8

3.2.4 Model Validation 9

3.3 Statistical Analysis of Waiting Times 10

3.4 Model A: The Exponential Distribution 11

3.4.1 The Exponential Model 11

3.4.2 Maximum Likelihood Estimation 11

3.4.3 Expected Waiting Time 12

3.5 Model B: The Weibull Distribution 13

3.5.1 The Weibull Model 13

3.5.2 Likelihood and Partial Derivatives 13

3.5.3 The Log-Likelihood Function 14

3.5.4 Maximising the Log-Likelihood 15

3.5.5 Expected Waiting Time 16

3.6 Comparing the Two Models 17

3.6.1 Graphical Comparison 17

3.6.2 Quantitative Comparison 18

3.6.3 Interpreting the Shape Parameter 19

3.7 Predicting the Next Record 20

3.7.1 Point Predictions 20

3.7.2 95% Prediction Intervals 21

3.7.3 Checking the Predictions Against Reality 21

4 Conclusion **22**

5 Evaluation **23**

5.1 Strengths and Weaknesses of the Investigation 23

5.2 Possible Extensions and Further Investigation 24

References **25**

Appendices **26**

1 Introduction and Rationale

A speedrun is an attempt to complete a video game, or part of a game, as quickly as possible (Cambridge Dictionary, 2024). Elite players around the world compete in speedrunning by repeatedly refining their strategies and techniques in an effort to surpass the world record.

Video games have been my greatest passion since I was eight years old. Over the years I played games such as *Minecraft*, *Roblox* and *Clash Royale* to near-mastery, and as a result I was naturally exposed to the speedrunning communities of several games. When I first attempted a speedrun, my record improved rapidly through small adjustments to my strategy. After hundreds of attempts, however, I could no longer find shortcuts or further refinements. My own record progression showed diminishing returns: large gains at the early stage, followed by small increments that required perfect execution. Curious whether other players experienced the same trend, I examined the global speedrun leaderboards. There I noticed a similar pattern: world record times dropped sharply at first and then plateaued as improvements became marginal. This suggested that record times might follow a predictable trend that could be described through statistical modelling, and it led me to ask whether I could predict when the next world record would be set, and what time it would achieve.

Although the exact date cannot be known in advance, I realised that the knowledge I had gained in the Mathematics: Analysis and Approaches course could be used to estimate both the date and the time of the next record. I have recently been speedrunning the physics-based climbing game *Getting Over It with Bennett Foddy* (GOIWBF). I chose this game not only because I am currently playing it, but also because it has a fairly large speedrunning community with a wealth of historical data.

While planning the investigation, I also realised that the answer depends heavily on how the waiting time between records is modelled. The simplest choice, the exponential distribution, assumes that records are broken at a constant rate. My own experience suggests otherwise: the longer a record stands, the harder it seems to break. I therefore decided to model the waiting times in two ways, with an exponential distribution and with a more flexible Weibull distribution, and to compare which better describes the speedrunning community and which produces the more realistic prediction.

2 Aim and Approach

The aim of this investigation is to model the progression of GOIWBF speedrun world records over time and to predict both the date and the time of the next world record.

I construct two complementary types of model: one for the record times themselves, to understand how they improve over time, and one for the waiting times between records, to explain the gaps between successive record breaks. The following models are developed.

Record progression model. Historical data show world record times decreasing non-linearly over the years: early records dropped dramatically, whereas later improvements slowed as record times approached the limits of skill. This model describes record time as a function of record date, capturing the diminishing rate of improvement observed both in the speedrunning community and in my own experience. I derive it by regression analysis using logarithmic linearisation and least-squares estimation.

Waiting time models. I examine histograms of the waiting times between successive records and then fit two candidate distributions by maximum likelihood estimation:

- *Model A: the exponential distribution*, which has one parameter and assumes a constant rate of record-breaking. Its estimator can be found with single-variable calculus.
- *Model B: the Weibull distribution*, which has two parameters and allows the rate of record-breaking to increase or decrease with the time since the last record. Its estimators require partial differentiation and a numerical solution.

I then compare the two models graphically and with quantitative measures of fit. For each model, I compute the expected waiting time from the definition of the expected value of a continuous random variable and a 95% prediction interval from the cumulative distribution function, and I convert these into a predicted record date and record time using the progression model. Finally, I check both predictions against what actually happened and discuss which model is more realistic.

All speedrun data were taken from the official `speedrun.com` leaderboard for GOIWBF to ensure that the records are verified. The dataset contains 32 verified world records with their exact dates and times. Intermediate calculations were carried out in Microsoft Excel using unrounded values, because Excel performs transparent, manually verifiable calculations to 15 significant figures. Where an equation could not be solved by hand, I used Python, and Desmos was used to visualise curves and probability density functions. Final results are reported to three decimal places for clarity, but unrounded values were used in all calculations.

3 Exploration

3.1 Definitions and Data Collection

I collected 32 verified world records for GOIWBF (glitchless category), spanning 2019 to 2025, from the official leaderboard on `speedrun.com` (Speedrun.com, 2017). On inspection, I noticed two patterns in this dataset: a consistent downward trend in record times, and clusters of records separated by gaps of irregular length.

Before performing any statistical analysis, I indexed the records chronologically by $i = 1, 2, \ldots, 32$. Let t_i be the date of the ith record measured in years since the first record, so that $t_1 = 0$ corresponds to the calendar year 2019.217. The sequence is therefore strictly increasing: $0 = t_1 < t_2 < \cdots < t_{32}$. Let R_i denote the corresponding record time in seconds. By definition, each new record is faster than the previous one, so $R_1 > R_2 > \cdots > R_{32}$. I then defined the waiting time between successive records as

$$x_i = t_{i+1} - t_i, \qquad i = 1, 2, \ldots, 31, \tag{1}$$

where each $x_i > 0$ is measured in years. This gives three data series:

$$\{t_1, t_2, \ldots, t_{32}\}, \qquad \{R_1, R_2, \ldots, R_{32}\}, \qquad \{x_1, x_2, \ldots, x_{31}\}. \tag{2}$$

The full dataset is given in Appendix 1 and the waiting times in Appendix 2.

3.2 Modelling Speedrun World Record Progression

To understand the trend in the data, I first produced a scatter plot of the record progression. Figure 1 plots the world record time R_i (seconds) against the record date t_i (years). Each of the 32 points represents one verified record.

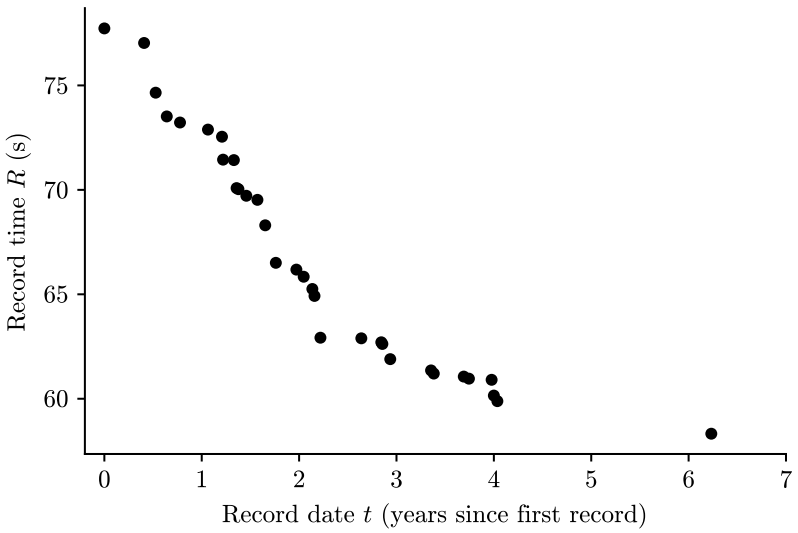


Figure 1: World record time R_i against record date t_i.

The plot shows a clear negative trend. Later records always have lower times because, by definition, a new record must be faster than the previous one. Qualitatively, the rate of improvement appears to slow over time: early records show larger drops (roughly 1 to 3 seconds), whereas recent records show only marginal improvements (roughly 0.03 to 0.7 seconds). This may indicate that the community is approaching the limits of human performance or of the game itself, so that new records become progressively harder to set. Because this pattern of diminishing improvement suggested that a non-linear model would be more appropriate than a linear one, I compared three models (linear, logarithmic and exponential) using Excel's trendline feature. Table 1 summarises the results.

Model	Function	R^2	Limitation
Linear	$y = -3.786t + 75.154$	0.865	Does not capture the pattern of rapid early improvement followed by slowing improvement.
Logarithmic	$y = -3.984 \ln t + 68.711$	varies with t_1	Undefined at $t_1 = 0$; predicts negative record times in the long term. (Shown for $t_1 = 0.01$, where $R^2 = 0.657$.)
Exponential	$y = 75.460\, e^{-0.057t}$	0.881	Asymptote at $y = 0$.

Table 1: Summary of the candidate progression models and their limitations.

The linear model gives a slightly poorer fit than the exponential model ($R^2 = 0.865$). More importantly, it cannot explain the rapid early improvements followed by the flattening of record times, so it would underestimate early improvements and overestimate recent ones. The logarithmic model describes this trend better visually, but it requires the first

data point to be adjusted arbitrarily (for example, to $t_1 = 0.01$ or $t_1 = 0.001$), because $\ln t$ is undefined at $t = 0$. As Figure 2 shows, shifting t_1 slightly away from zero produces very different values of R^2. The logarithmic model also predicts physically impossible negative record times in the long term (as $t \to \infty$, $y \to -\infty$), which makes it unsuitable for forecasting.

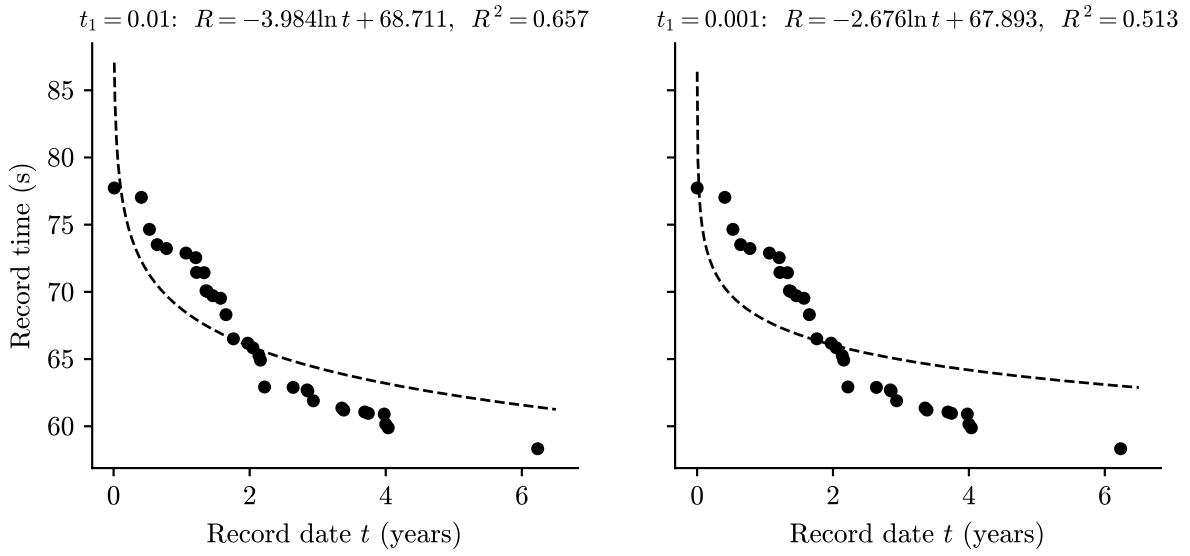


Figure 2: Logarithmic model fitted with two different values of t_1.

The exponential model avoids the limitations of both alternatives. It fits the data well without redefining t_1, and it captures diminishing returns, in which each optimisation of strategy yields a smaller gain in time as the optimal routes become saturated. It also reproduces the large improvements of the early phase and an asymptotic approach towards a skill ceiling, without the record time ever becoming negative. This mirrors my own practice sessions, in which my first optimisations saved seconds but later adjustments saved only milliseconds. One limitation is that $y \to 0$ as $t \to \infty$, whereas completing the game in zero seconds is impossible. Nevertheless, the exponential model combines empirical fit and physical realism better than the linear and logarithmic models, so I chose it to model the progression of GOIWBF records.

3.2.1 Exponential Regression Model

Let $T(t)$ denote the predicted record time at date t (years). The exponential model takes the form

$$T(t) = a\, e^{bt}, \qquad a > 0, \quad b < 0, \tag{3}$$

where $a = T(0)$ is the modelled record time at the first record and b controls the rate of improvement. Since $b < 0$, $T(t)$ decreases at a diminishing rate as t increases.

3.2.2 Logarithmic Linearisation

Since fitting a non-linear model directly is complex, I took the natural logarithm of both sides to linearise the model:

$$\ln T(t) = \ln a + bt. \tag{4}$$

Letting $y_i = \ln R_i$ and $\alpha = \ln a$, and assuming $T(t_i) \approx R_i$, the model becomes

$$y_i = \alpha + bt_i + \varepsilon_i, \tag{5}$$

where ε_i represents random error. Since the pairs (t_i, y_i) follow an approximately linear relationship, α and b can be estimated by least-squares linear regression.

3.2.3 Least-Squares Estimation

For the linear model $y = \alpha + bt$, the least-squares estimates are

$$b = \frac{S_{ty}}{S_{tt}}, \qquad \alpha = \bar{y} - b\bar{t}, \tag{6}$$

where

$$\bar{t} = \frac{1}{32}\sum_{i=1}^{32} t_i, \qquad \bar{y} = \frac{1}{32}\sum_{i=1}^{32} y_i,$$

$$S_{tt} = \sum_{i=1}^{32}(t_i - \bar{t})^2, \qquad S_{yy} = \sum_{i=1}^{32}(y_i - \bar{y})^2, \qquad S_{ty} = \sum_{i=1}^{32}(t_i - \bar{t})(y_i - \bar{y}). \tag{7}$$

Table 2 gives these statistics for the GOIWBF dataset.

Statistic	Value
$\bar{t}$	2.206
$\bar{y}$	4.198
S_{tt}	57.157
S_{yy}	0.209
S_{ty}	−3.243

Table 2: Least-squares summary statistics for the linearised model.

The negative value of S_{ty} confirms the negative relationship between log-record time and date. The least-squares estimates are

$$b = \frac{S_{ty}}{S_{tt}} \approx -0.057, \qquad \alpha = \bar{y} - b\bar{t} \approx 4.324, \qquad a = e^{\alpha} \approx 75.460, \tag{8}$$

which give the fitted progression model

$$T(t) = 75.460\, e^{-0.057t}. \tag{9}$$

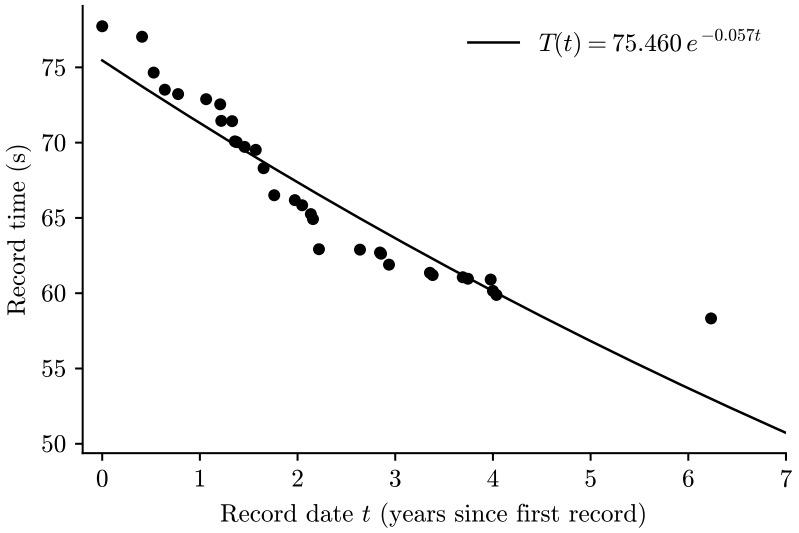


Figure 3: Fitted exponential regression model overlaid on the scatter plot.

Figure 3 shows the model overlaid on the data. Although it is the best fit in the least-squares sense, it does not describe the data perfectly. Its curvature is lower than I expected: it underestimates R_1 and R_2 and overestimates R_{20}, R_{21} and R_{24} by noticeably more than the other records. In addition, the 32nd record lies well above the curve, which indicates that recent improvement has been slower than the model predicts, so the model is likely to be over-optimistic about future gains. Since $T(t) \to 0$ as $t \to \infty$, the model should be used only for short- to medium-term forecasts. As this investigation predicts only the next record, the limitations of long-term extrapolation have little effect on the results.

3.2.4 Model Validation

The Pearson correlation coefficient for the linearised data is

$$r = \frac{S_{ty}}{\sqrt{S_{tt} S_{yy}}} = \frac{-3.243}{\sqrt{57.157 \times 0.209}} \approx -0.939. \tag{10}$$

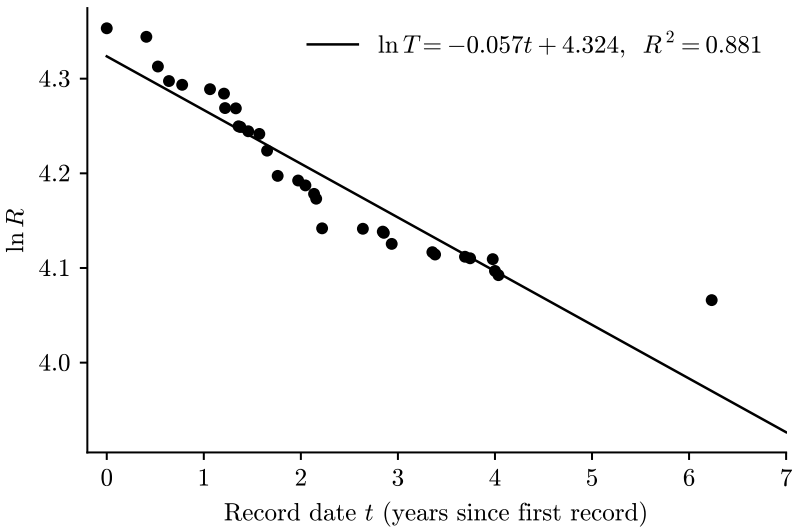


Figure 4: Fitted linear regression on the log-transformed data ($y_i = \ln R_i$ against t_i).

Figure 4 and the value $|r| = 0.939$ show a strong linear relationship between t_i and $\ln R_i$. The coefficient of determination, $R^2 \approx 0.881$, indicates that the model explains 88.1% of the variance in $\ln R_i$, so I proceeded with this progression model for the predictions that follow.

3.3 Statistical Analysis of Waiting Times

To predict how long it might take for the current record to be broken, I modelled the historical waiting times between records. I began by plotting histograms of the 31 waiting times $X = \{x_1, x_2, \ldots, x_{31}\}$, since these show the shape, spread and skewness of the data and any possible multimodality.

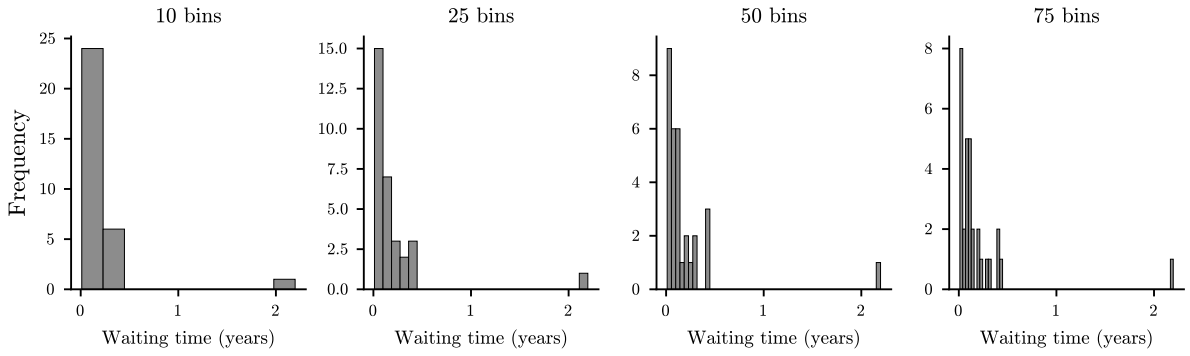


Figure 5: Histograms of the waiting times x_i with 10, 25, 50 and 75 bins.

In each histogram in Figure 5, the distribution is strongly right-skewed. Most waiting times are short (20 of the 31 are under 0.12 years, or about six weeks), but there is a long

right tail, including one exceptional gap of about 2.2 years before the most recent record. This shape is consistent with a Poisson-like process in which records occur randomly. It also reflects two features of speedrunning that I have observed in the community: bursts of records in quick succession after a major discovery, such as a new route or glitch that triggers a series of follow-up optimisations, and long periods without records when no new ideas emerge.

These observations suggested two candidate distributions. The exponential distribution is the natural starting point because it describes the time between events in a Poisson process, but it assumes that the chance of a new record in any short interval stays the same no matter how long the current record has stood. Survival analysis (time-to-event analysis), which studies the time until an event occurs (UCL Cancer Trials Centre, n.d.), offers a more flexible alternative: the Weibull distribution, which allows this chance to change over time (Meeker & Escobar, 1998). Since my experience suggests that a record becomes harder to break the longer it survives, I fitted both distributions and compared them.

3.4 Model A: The Exponential Distribution

3.4.1 The Exponential Model

The exponential distribution is a continuous probability distribution that models the time elapsed before an event occurs. It is often used to answer questions such as "How long will a machine work before it breaks down?". Using it requires the assumption that waiting times are independent, which may not hold perfectly given the dynamics of speedrunning, but it is consistent with the requirement that waiting times be positive and resembles the histograms. Under this model, $X \sim \text{Exp}(\lambda)$, where λ is the rate of record-breaking (records per year), with probability density function (pdf) (Taboga, 2021)

$$f_X(x) = \begin{cases} \lambda e^{-\lambda x}, & x \geq 0, \\ 0, & x < 0. \end{cases} \tag{11}$$

3.4.2 Maximum Likelihood Estimation

To find the value of λ that makes the 31 observed waiting times most probable, I used maximum likelihood estimation. The likelihood function $L(\lambda \mid X)$ is the product of the individual pdfs evaluated at each observation, regarded as a function of λ for a fixed set of observations (Watkins, 2011). For example, the likelihood of observing both x_1 and x_2 is

$$L(\lambda \mid x_1, x_2) = f_X(x_1 \mid \lambda) \cdot f_X(x_2 \mid \lambda) = \lambda e^{-\lambda x_1} \cdot \lambda e^{-\lambda x_2} = \lambda^2 e^{-\lambda(x_1+x_2)}. \tag{12}$$

Assuming that the waiting times are independent and identically distributed, the likelihood of the entire dataset is

$$L(\lambda \mid X) = \prod_{i=1}^{31} \lambda e^{-\lambda x_i} = \lambda^{31} e^{-\lambda \sum_{i=1}^{31} x_i}. \tag{13}$$

A large value of λ makes λ^{31} large but $e^{-\lambda \sum x_i}$ very small, and a small value has the opposite effect, so there is a balance point that maximises $L(\lambda)$. Before differentiating, I took the logarithm of $L(\lambda)$, which converts the product into a sum. Since the logarithm is an increasing function, maximising $L(\lambda)$ is equivalent to maximising the log-likelihood (Appendix 3):

$$\ell(\lambda) = \ln L(\lambda) = 31 \ln \lambda - \lambda \sum_{i=1}^{31} x_i. \tag{14}$$

Setting the derivative equal to zero,

$$\frac{d\ell}{d\lambda} = \frac{31}{\lambda} - \sum_{i=1}^{31} x_i = 0 \quad \Longrightarrow \quad \hat{\lambda} = \frac{31}{\sum_{i=1}^{31} x_i} = \frac{31}{6.233} \approx 4.974 \text{ records per year.} \tag{15}$$

The second derivative test confirms that this is a maximum, since

$$\frac{d^2\ell}{d\lambda^2} = -\frac{31}{\lambda^2} \approx -1.253 < 0. \tag{16}$$

The fitted pdf is therefore

$$f_X(x) = 4.974\, e^{-4.974x}, \qquad x \geq 0. \tag{17}$$

3.4.3 Expected Waiting Time

The expected waiting time is $\mathbb{E}[X] = \int_{-\infty}^{\infty} x f_X(x)\, dx = \int_0^{\infty} x \cdot 4.974\, e^{-4.974x}\, dx$, since $f_X(x) = 0$ for $x < 0$. Using integration by parts with

$$u = x, \qquad dv = 4.974\, e^{-4.974x}\, dx \quad \Longrightarrow \quad du = dx, \qquad v = -e^{-4.974x}, \tag{18}$$

I obtained (full working in Appendix 4)

$$\mathbb{E}[X] = \Big[-xe^{-4.974x} \Big]_0^{\infty} + \int_0^{\infty} e^{-4.974x}\, dx = \lim_{x \to \infty} \left(-\frac{x}{e^{4.974x}} \right) + \frac{1}{4.974}. \tag{19}$$

The limit has the indeterminate form $-\infty/\infty$, so by L'Hôpital's rule

$$\lim_{x \to \infty} \left(-\frac{x}{e^{4.974x}} \right) = \lim_{x \to \infty} \left(-\frac{1}{4.974\, e^{4.974x}} \right) = 0, \tag{20}$$

and therefore

$$\mathbb{E}[X] = \frac{1}{4.974} \approx 0.201 \text{ years} \approx 73.4 \text{ days}. \tag{21}$$

3.5 Model B: The Weibull Distribution

3.5.1 The Weibull Model

The two-parameter Weibull distribution is a continuous probability distribution characterised by two positive parameters (ScienceDirect, n.d.). Its pdf is (Meeker & Escobar, 1998)

$$f(x;\eta,\beta) = \frac{\beta}{\eta}\left(\frac{x}{\eta}\right)^{\beta-1} e^{-(x/\eta)^{\beta}}, \qquad x \geq 0, \tag{22}$$

where $\eta > 0$ is the *scale* parameter, which controls the spread of the distribution, and $\beta > 0$ is the *shape* parameter, which determines its shape. What makes the Weibull distribution useful here is that when $\beta = 1$, the pdf becomes $\frac{1}{\eta}e^{-x/\eta}$, which is exactly the exponential distribution with $\lambda = 1/\eta$. The exponential model is therefore a special case of the Weibull model, and the estimated value of β will tell me whether the extra flexibility is needed. As I show in Section 3.6, $\beta < 1$ corresponds to a rate of record-breaking that decreases the longer a record stands, and $\beta > 1$ to one that increases.

3.5.2 Likelihood and Partial Derivatives

As before, assuming that the waiting times are independent, the likelihood function is the product of the pdfs at each observation, but it is now a function of two parameters:

$$L(\eta,\beta) = \prod_{i=1}^{31} f(x_i \mid \eta,\beta) = \prod_{i=1}^{31} \frac{\beta}{\eta}\left(\frac{x_i}{\eta}\right)^{\beta-1} e^{-(x_i/\eta)^{\beta}}. \tag{23}$$

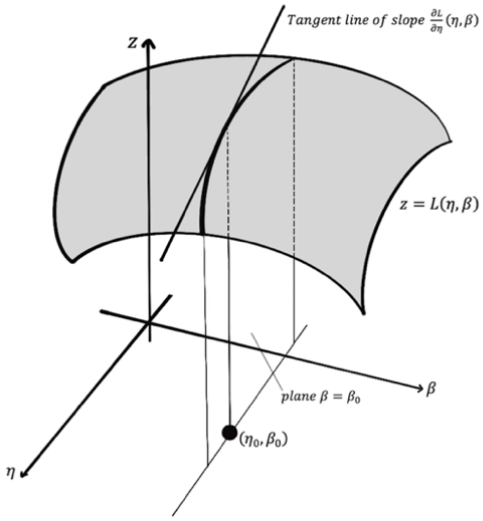


Figure 6: Visual representation of a likelihood surface $z = L(\eta, \beta)$ and the slope of a tangent in the η-direction.

To maximise a function of two variables, it helps to visualise $z = L(\eta, \beta)$ as a surface above the (η, β)-plane (Figure 6). A partial derivative is the instantaneous rate of change of the function with respect to one variable while the other is held fixed. For example, moving through a point (η_0, β_0) parallel to the η-axis, β remains fixed at β_0, and differentiating with respect to η gives the partial derivative $\partial L/\partial \eta$: the slope of the surface in the η-direction (East Tennessee State University, n.d.). At the highest point of the surface, the slope must be zero in both directions, so the maximum likelihood estimates $\hat{\eta}$ and $\hat{\beta}$ are found by setting both partial derivatives equal to zero simultaneously (May & Bart, n.d.).

3.5.3 The Log-Likelihood Function

As for the exponential model, I took logarithms to turn the product into a sum (Taboga, 2019). Expanding the logarithm of each term (full working in Appendix 5),

$$\begin{aligned} \ell(\eta, \beta) = \ln L(\eta, \beta) &= \sum_{i=1}^{n} \left[\ln \beta - \ln \eta + (\beta - 1)(\ln x_i - \ln \eta) - \left(\frac{x_i}{\eta} \right)^{\beta} \right] \\ &= n \ln \beta - n\beta \ln \eta + (\beta - 1) \sum_{i=1}^{n} \ln x_i - \sum_{i=1}^{n} \left(\frac{x_i}{\eta} \right)^{\beta}, \end{aligned} \tag{24}$$

where $n = 31$.

3.5.4 Maximising the Log-Likelihood

Partial derivative with respect to η. Since the derivative of a sum is the sum of the derivatives (proved in Appendix 6), and x_i^β is constant with respect to η,

$$\frac{\partial \ell}{\partial \eta} = -\frac{n\beta}{\eta} - \sum_{i=1}^{n} x_i^\beta \cdot \frac{\partial}{\partial \eta}\left(\eta^{-\beta}\right) = -\frac{n\beta}{\eta} + \frac{\beta}{\eta^{\beta+1}} \sum_{i=1}^{n} x_i^\beta. \tag{25}$$

Setting this equal to zero and multiplying by $\eta^{\beta+1}/\beta$ gives $\sum x_i^\beta - n\eta^\beta = 0$, so

$$\eta = \left(\frac{1}{n} \sum_{i=1}^{n} x_i^\beta\right)^{1/\beta}. \tag{26}$$

This expresses the best value of η in terms of β, which reduces the problem to finding a single unknown.

Partial derivative with respect to β. Using $\frac{d}{d\beta} A^\beta = A^\beta \ln A$ with $A = x_i/\eta$ (full working in Appendix 7),

$$\frac{\partial \ell}{\partial \beta} = \frac{n}{\beta} - n \ln \eta + \sum_{i=1}^{n} \ln x_i - \sum_{i=1}^{n} \left(\frac{x_i}{\eta}\right)^\beta \ln \frac{x_i}{\eta}. \tag{27}$$

Substituting $\eta^\beta = \frac{1}{n} \sum x_i^\beta$ from (26), the terms involving $\ln \eta$ cancel, leaving an equation in β alone:

$$g(\beta) = \frac{n}{\beta} + \sum_{i=1}^{n} \ln x_i - \frac{n \sum_{i=1}^{n} x_i^\beta \ln x_i}{\sum_{i=1}^{n} x_i^\beta} = 0. \tag{28}$$

Solving numerically. Equation (28) is transcendental, because β appears both outside and inside exponents, so it cannot be solved algebraically. With $n = 31$ and $\sum \ln x_i \approx -73.033$, I solved it numerically in Python using a root-finding method. Figure 7 shows that $g(\beta)$ is decreasing and crosses zero exactly once for $\beta > 0$, at

$$\hat{\beta} \approx 0.806. \tag{29}$$

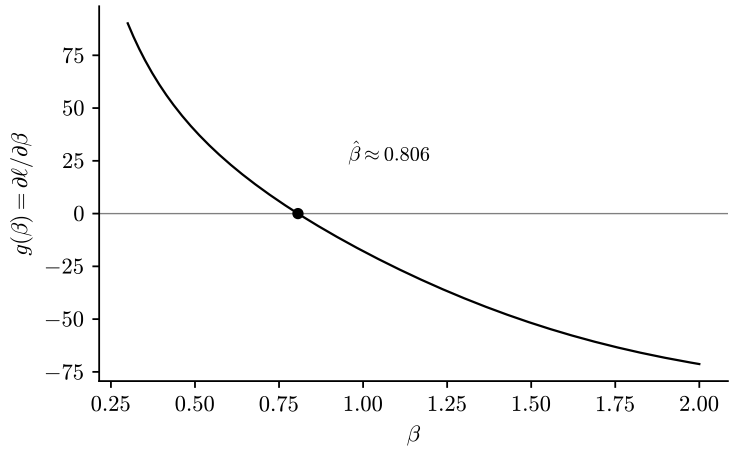


Figure 7: The function $g(\beta) = \partial\ell/\partial\beta$ after substituting (26); its root gives $\hat{\beta}$.

Substituting $\hat{\beta}$ into (26) gives

$$\hat{\eta} = \left(\frac{1}{31}\sum_{i=1}^{31} x_i^{0.806}\right)^{1/0.806} \approx 0.172 \text{ years.} \tag{30}$$

To confirm that this critical point is a maximum rather than a saddle point, I computed the second partial derivatives of ℓ numerically at $(\hat{\eta}, \hat{\beta})$:

$$\frac{\partial^2\ell}{\partial\eta^2} \approx -683.04, \qquad \frac{\partial^2\ell}{\partial\beta^2} \approx -113.85, \qquad \frac{\partial^2\ell}{\partial\eta\,\partial\beta} \approx 93.999. \tag{31}$$

The determinant of the Hessian matrix is $(-683.04)(-113.85) - (93.999)^2 \approx 68\,926 > 0$, and $\partial^2\ell/\partial\eta^2 < 0$, so by the second derivative test for functions of two variables, $(\hat{\eta}, \hat{\beta})$ is a local maximum. As a further check, the estimates agree with those from the built-in fitting routine of the SciPy library to three decimal places. Substituting the estimates into (22) gives the fitted Weibull pdf

$$f(x) = 3.337\, x^{-0.194}\, e^{-4.139\, x^{0.806}}, \qquad x \geq 0. \tag{32}$$

3.5.5 Expected Waiting Time

Substituting the Weibull pdf into the definition of the expected value and using the substitution $u = (x/\eta)^\beta$ (full working in Appendix 8) gives

$$\mathbb{E}[X] = \int_0^\infty x\, f(x)\, dx = \eta \int_0^\infty u^{1/\beta} e^{-u}\, du. \tag{33}$$

The remaining integral is the gamma function, $\Gamma(\alpha) = \int_0^\infty u^{\alpha-1}e^{-u}\,du$ (Frost, 2021), with $\alpha - 1 = 1/\beta$. Hence

$$\mathbb{E}[X] = \eta\,\Gamma\left(1 + \frac{1}{\beta}\right) = 0.172 \times \Gamma(2.240) \approx 0.172 \times 1.127 \approx 0.194 \text{ years} \approx 70.7 \text{ days.} \quad (34)$$

As a check, setting $\beta = 1$ gives $\mathbb{E}[X] = \eta\,\Gamma(2) = \eta = 1/\lambda$, which agrees with the exponential result.

3.6 Comparing the Two Models

3.6.1 Graphical Comparison

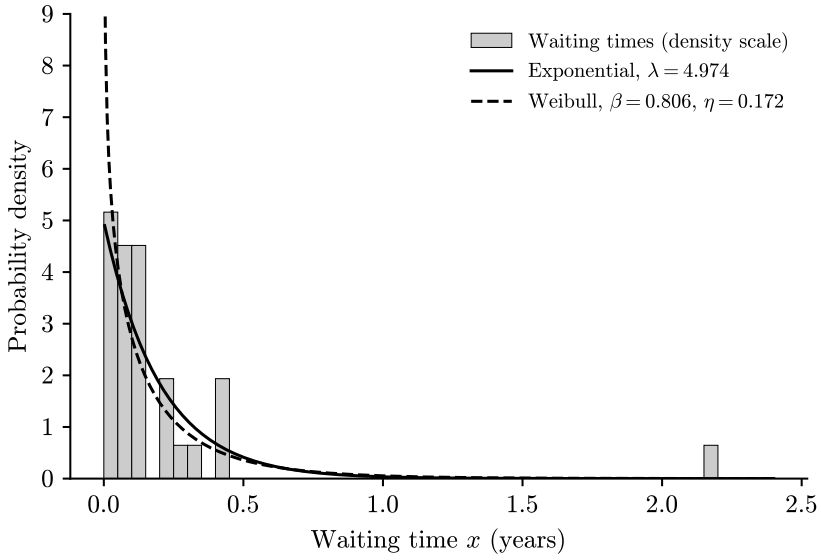


Figure 8: Fitted exponential and Weibull pdfs overlaid on a histogram of the waiting times, scaled to unit area.

Figure 8 shows both fitted pdfs overlaid on a histogram of the waiting times, scaled so that the total area of the bars is 1 and the histogram is directly comparable with the pdfs. Both curves follow the general shape of the data, but they differ at the two extremes. The Weibull pdf is much higher near $x = 0$ (it tends to infinity as $x \to 0$, because $\beta < 1$), which matches the large number of very short gaps produced by bursts of records. It also decays more slowly in the right tail, which better accommodates long gaps.

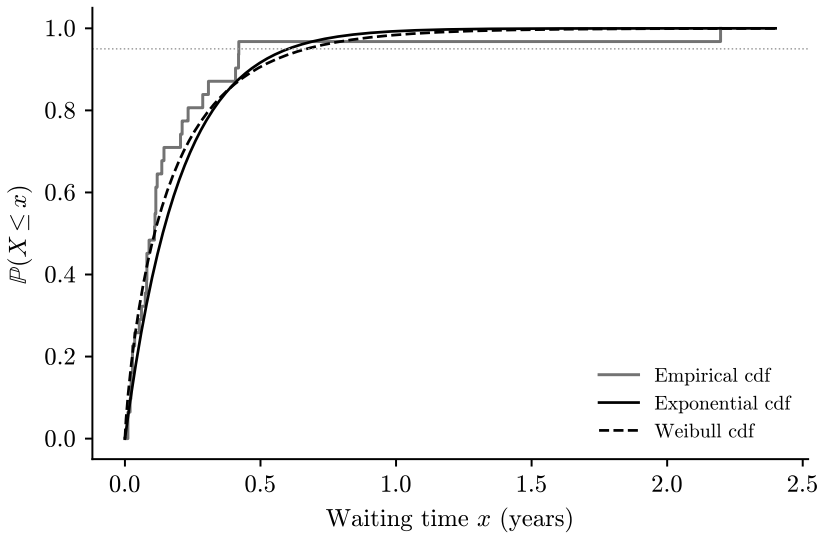


Figure 9: Empirical cumulative distribution function of the waiting times compared with the fitted exponential and Weibull cdfs. The dotted line marks a cumulative probability of 0.95.

Figure 9 compares the cumulative distribution functions (cdfs) with the empirical cdf, which rises by 1/31 at each observed waiting time. The Weibull cdf follows the empirical cdf more closely, particularly for short waiting times, where the exponential cdf rises too slowly.

3.6.2 Quantitative Comparison

To compare the models more objectively, I used two measures, summarised in Table 3.

1. **Maximised log-likelihood.** The higher the value of ℓ at the estimates, the more probable the observed data are under the model. The Weibull model has the higher value, but this is guaranteed, because it contains the exponential model as a special case and has an extra parameter.

2. **Likelihood ratio test.** Since the Weibull model always fits at least as well, the real question is whether its improvement is large enough to be genuine, rather than a side effect of having an extra parameter. Because the exponential model is the Weibull model with $\beta = 1$, this is equivalent to testing

$$H_0 : \ \beta = 1 \text{ (the exponential model is adequate)} \qquad \text{against} \qquad H_1 : \ \beta \neq 1.$$

   The test statistic is twice the difference between the maximised log-likelihoods,

$$2\Big(\ell_{\text{Weibull}} - \ell_{\text{Exp}}\Big) = 2(20.504 - 18.728) \approx 3.55.$$

If H_0 is true, this statistic approximately follows a χ^2 distribution with one degree of freedom, because the Weibull model has one extra parameter (Wilks, 1938). At the 5% significance level, the critical value is $\chi^2_1(0.05) \approx 3.841$. Since $3.55 < 3.841$ (equivalently, $p \approx 0.059 > 0.05$), I do not reject H_0. The evidence that $\beta < 1$ is therefore suggestive but not conclusive, which is unsurprising with only 31 observations.

	Exponential	Weibull
Parameters	$\hat{\lambda} = 4.974$	$\hat{\beta} = 0.806$, $\hat{\eta} = 0.172$
Number of parameters k	1	2
Maximised log-likelihood ℓ	18.728	20.504
Likelihood ratio statistic	3.551 ($p \approx 0.059$)	
$\mathbb{P}(X \geq 2.197)$	1.8×10^{-5}	4.1×10^{-4}

Table 3: Comparison of the fitted exponential and Weibull models.

The last row of Table 3 is particularly revealing. Using the survival function $\mathbb{P}(X \geq x) = 1 - F(x)$, the exponential model gives the observed 2.197-year gap before the 32nd record a probability of only about $e^{-4.974 \times 2.197} \approx 1.8 \times 10^{-5}$, whereas the Weibull model gives $e^{-(2.197/0.172)^{0.806}} \approx 4.1 \times 10^{-4}$, which is more than 20 times larger. Neither model expects such a long gap, but the exponential model treats it as almost impossible.

3.6.3 Interpreting the Shape Parameter

The difference between the models can be understood through the *hazard rate* $h(x) = f(x)/(1 - F(x))$, which gives the instantaneous rate at which a record is broken, given that it has survived for time x. For the exponential model,

$$h(x) = \frac{\lambda e^{-\lambda x}}{e^{-\lambda x}} = \lambda, \tag{35}$$

which is constant: the model assumes that a record that has stood for two years is as easy to break as one set yesterday. This is the memoryless property. For the Weibull model,

$$h(x) = \frac{\frac{\beta}{\eta}\left(\frac{x}{\eta}\right)^{\beta-1} e^{-(x/\eta)^{\beta}}}{e^{-(x/\eta)^{\beta}}} = \frac{\beta}{\eta}\left(\frac{x}{\eta}\right)^{\beta-1}, \tag{36}$$

which is a decreasing function of x whenever $\beta < 1$, because the exponent $\beta - 1$ is negative.

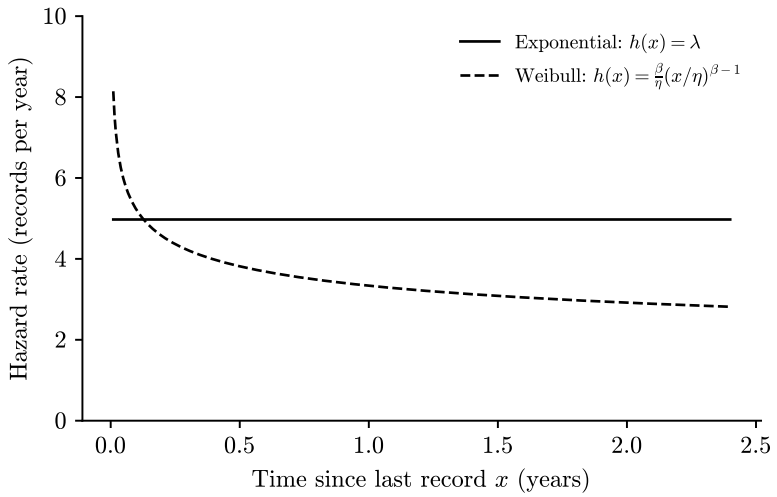


Figure 10: Hazard rates of the fitted models. The Weibull hazard decreases as a record ages.

With $\hat{\beta} \approx 0.806$, the Weibull hazard falls from about 6.0 records per year when a record is two to three weeks old to about 3.7 after seven months and 2.9 after two years (Figure 10). In other words, the fitted Weibull model says that the longer a record survives, the harder it becomes to break. This matches what I experienced as a speedrunner: a record that survives the first wave of attempts is usually one that has exhausted the easy improvements, and beating it requires a genuinely new idea. It is satisfying that this intuition emerged from the data rather than being imposed on the model.

3.7 Predicting the Next Record

3.7.1 Point Predictions

The next record date is the date of the current (32nd) record plus a random waiting time:

$$t_{\text{next}} = t_{32} + X, \qquad t_{32} = 6.233 \text{ (14 June 2025)}. \tag{37}$$

Adding each model's expected waiting time gives the expected record date, and evaluating the progression model (9) at that date gives the expected record time.

For the exponential model,

$$t_{\text{exp}} = 6.233 + 0.201 = 6.434 \text{ (26 August 2025)}, \qquad T(6.434) \approx 52.383 \text{ s}. \tag{38}$$

For the Weibull model,

$$t_{\text{exp}} = 6.233 + 0.194 = 6.427 \text{ (24 August 2025)}, \qquad T(6.427) \approx 52.405 \text{ s}. \tag{39}$$

The two point predictions are almost identical, because both models were fitted to the same data and have similar means. The models differ much more in how they spread the probability around this mean, which becomes clear in the prediction intervals.

3.7.2 95% Prediction Intervals

For each model, I constructed a 95% prediction interval for the waiting time of the form $(0, x)$, where $\mathbb{P}(X \leq x) = 0.95$. The cdf of the exponential model is obtained by integrating its pdf:

$$F_X(x) = \int_0^x 4.974\, e^{-4.974s}\, ds = 1 - e^{-4.974x}, \tag{40}$$

so (Appendix 9)

$$1 - e^{-4.974x} = 0.95 \quad \Longrightarrow \quad x = -\frac{\ln 0.05}{4.974} \approx 0.602 \text{ years } (\approx 220 \text{ days}). \tag{41}$$

For the Weibull model, the cdf is $F(x) = 1 - e^{-(x/\eta)^\beta}$ (Appendix 9), so

$$1 - e^{-(x/\eta)^\beta} = 0.95 \quad \Longrightarrow \quad x = \eta\,(-\ln 0.05)^{1/\beta} \approx 0.670 \text{ years } (\approx 245 \text{ days}). \tag{42}$$

The Weibull interval is wider because its heavier tail assigns more probability to long waits. Table 4 converts both intervals into dates and record times.

Model	Prediction	Date	Record time (s)
Exponential	Point prediction	26 August 2025	52.383
	95% prediction interval	14 June 2025 – 20 January 2026	51.204 – 52.984
Weibull	Point prediction	24 August 2025	52.405
	95% prediction interval	14 June 2025 – 14 February 2026	51.008 – 52.984

Table 4: Predictions for the next GOIWBF world record from each waiting time model.

3.7.3 Checking the Predictions Against Reality

At the time of writing, the upper bounds of both prediction intervals have passed without a new world record being set. Neither 95% interval contained the actual waiting time, so both predictions were unsuccessful. A 95% interval is expected to miss 5% of the time, so one miss does not by itself invalidate a model, but the comparison above suggests that the miss is not just bad luck. The exponential model is built on a constant rate, and the observed data, especially the 2.2-year gap, show that the rate has not been constant. The Weibull model moves in the right direction: it produced a later upper bound and gave a long gap a much higher probability. However, it still underestimated the wait, because a

single value of β cannot capture how much the community's activity has slowed in recent years.

The two models also differ in how they would update a prediction as time passes, which is where the memoryless property matters most. Suppose that on 20 January 2026, the upper bound of the exponential interval, I had wanted to update my forecast given that the record was still standing. Because the exponential model is memoryless, it would simply predict a further average wait of $1/\lambda \approx 73$ days, that is, around 3 April 2026, exactly as if the record had just been set. The Weibull model instead uses the information that the record has survived. Its expected remaining waiting time, given that the record has already lasted s years, is

$$\mathbb{E}[X - s \mid X > s] = \frac{\int_s^\infty \big(1 - F(x)\big)\, dx}{1 - F(s)}, \tag{43}$$

which, evaluated numerically at $s = 0.602$, gives about 0.291 years (106 days), or around 6 May 2026. Because its hazard rate decreases, the Weibull model becomes more patient the longer a record survives, which is exactly the behaviour I would expect from an experienced speedrunner watching a strong record.

4 Conclusion

In this investigation, I modelled the record progression of *Getting Over It with Bennett Foddy* and predicted the date and time of the next (33rd) world record, using two different models for the waiting time between records. The progression model $T(t) = 75.460\, e^{-0.057t}$, obtained by logarithmic linearisation and least squares, explained 88.1% of the variance in log-record time and captured the diminishing returns of speedrunning. For the waiting times, the exponential model, fitted by single-variable maximum likelihood, predicted the next record on 26 August 2025 at 52.383 seconds, with a 95% prediction interval ending on 20 January 2026. The Weibull model, fitted by partial differentiation and a numerical solution, predicted almost the same date (24 August 2025, 52.405 seconds) but a wider interval ending on 14 February 2026.

Comparing the two models was the most valuable part of the investigation. The Weibull model achieved a higher likelihood, followed the histogram and the empirical cdf more closely, and gave the long 2.2-year gap a far more realistic probability. Its estimated shape parameter $\hat{\beta} \approx 0.806 < 1$ revealed a decreasing hazard rate: the longer a record stands, the less likely it is to be broken in any given week. This turned an intuition from my own speedrunning into a measurable property of the data. The likelihood ratio test showed that this evidence is not yet conclusive ($p \approx 0.059$), so with only 31 observations I cannot

rule out the simpler exponential model. Nevertheless, both models' intervals have passed without a new record, which confirms that the waiting time was underestimated. The main reason is that both models assume the waiting times are independent and come from a single, unchanging distribution, whereas the GOIWBF community has clearly moved into a slower, late phase in which improvements are much harder to find.

Through this investigation, I explored connections between several branches of mathematics, including differential and integral calculus, partial differentiation, numerical methods, probability and statistics. I was struck by how versatile the logarithm is: it linearised the progression model and turned both likelihood functions into sums that could be differentiated. Extending the exponential model to the Weibull model also showed me that a more general model can contain a simpler one as a special case, and that comparing them can reveal something the simpler model hides. Rather than simply enjoying speedrunning, I came to appreciate how mathematical tools can turn what feels like a random guess about the future into a quantitative prediction, while also exposing the assumptions behind it.

5 Evaluation

5.1 Strengths and Weaknesses of the Investigation

One strength of the progression model is its good fit to the historical data: it explains 88.1% of the variance in log-record time ($r = -0.939$). However, it assumes that record times improve at a constant relative rate b, which oversimplifies a process driven by players. When a meta shift occurs (a significant change in the most effective glitches, routes or strategies), improvements suddenly accelerate; this is followed by a period of stagnation, and the cycle tends to repeat. A single rate b cannot capture these non-stationary trends. The model also ignores physical limits: since $T(t) \to 0$ as $t \to \infty$, it would eventually predict impossibly short completion times, so it is reliable only in the short term. A model with a non-zero asymptote, such as $T(t) = c + a\,e^{bt}$, would be more realistic but would require non-linear fitting methods.

A major strength of the investigation is that the waiting times were modelled in two ways and compared using both graphs and quantitative measures. This showed not only which model fits better, but also why: the Weibull model's decreasing hazard rate captures the clustering of records followed by long droughts. Another strength is that the fitted pdfs matched the histograms well across different numbers of bins. However, both models share important weaknesses. First, both assume that waiting times are independent and identically distributed. In reality, a single strong player, a community effort, or new tools such as tool-assisted speedrun (TAS) analysis can produce clusters of records, so

consecutive waiting times are probably dependent. Second, both distributions are fixed over the whole history of the game. Even the Weibull model, whose hazard decreases with the age of each record, assumes that every record is drawn from the same distribution, whereas the community's overall activity has declined over time. Third, the parameters were estimated from only 31 waiting times. The single 2.197-year gap has a large influence on the estimates, and the small sample is why the likelihood ratio test could not distinguish clearly between the models. Finally, the Weibull estimates required a numerical solution, which I could verify graphically and with a second software package but not derive by hand.

These limitations made me ask whether it is reasonable to expect accurate predictions from a system driven by people. Speedrunning involves countless variables, such as player motivation, new discoveries and slumps in form, and modelling all of them is extremely difficult. From this, I learned that although future events are hard to predict precisely, mathematics can provide estimates, measure their uncertainty and show which assumptions matter most.

5.2 Possible Extensions and Further Investigation

One extension would address the assumption of a single distribution for the whole history of the game. I could divide the GOIWBF timeline into phases that match my experience as a speedrunner: an early phase (2019–2021, with frequent records from relatively easy improvements), a middle phase (2022–2023) and a late phase (2024 onwards, with rare records as optimisation saturates). Fitting a separate exponential or Weibull distribution to each phase would allow the rate of record-breaking to change over time, which would probably reduce the underestimation of the latest waiting time. I could also include indicator variables for known meta shifts so that these qualitative events enter the model.

A second extension would test whether my findings generalise to other games. In an earlier exploration, I fitted Weibull distributions to the waiting times between records for *Super Mario Bros.* (Any%) and *Minecraft*, and also obtained shape parameters below 1 ($\hat{\beta} \approx 0.911$ and 0.837 respectively). This suggests that a decreasing rate of record-breaking may be a common feature of speedrunning rather than something specific to GOIWBF. I would like to extend this comparison to games with major cycles of discovery or glitch-based categories, such as *The Legend of Zelda: Ocarina of Time*, where I would expect strongly clustered records followed by long periods of stagnation. Such a comparison could deepen our understanding of how human performance is optimised in competitive gaming.

References

Cambridge Dictionary. (2024, October 23). Speedrun. In *Cambridge Dictionary*. https://dictionary.cambridge.org/dictionary/english/speedrun

East Tennessee State University. (n.d.). *Functions of two variables*. Retrieved July 17, 2025, from https://math.etsu.edu/multicalc/prealpha/Chap2/Chap2-1/printversion.pdf

Frost, J. (2021, August 20). Gamma distribution: Uses, parameters & examples. *Statistics by Jim*. https://statisticsbyjim.com/probability/gamma-distribution/

May, M., & Bart, A. (n.d.). Critical points and extrema. In *Business calculus with Excel*. Saint Louis University. https://mathstat.slu.edu/~may/ExcelCalculus/sec-6-3-CriticalPointsExtrema.html

Meeker, W. Q., & Escobar, L. A. (1998). *Statistical methods for reliability data* (pp. 86–99). Wiley.

ScienceDirect. (n.d.). *Weibull distribution: An overview*. Retrieved July 17, 2025, from https://www.sciencedirect.com/topics/physics-and-astronomy/weibull-distribution

Speedrun.com. (2017). *Getting Over It with Bennett Foddy* [Leaderboard]. https://www.speedrun.com/goiwbf

Taboga, M. (2019). Log-likelihood. *StatLect*. https://www.statlect.com/glossary/log-likelihood

Taboga, M. (2021). Exponential distribution. In *Lectures on probability theory and mathematical statistics* (Online appendix). Kindle Direct Publishing. https://www.statlect.com/probability-distributions/exponential-distribution

UCL Cancer Trials Centre. (n.d.). *Outcome measures based on time-to-event data* [Training slides]. Retrieved July 17, 2025, from https://www.ctc.ucl.ac.uk/TrainingDocuments/Uploaded/Day%201%20-%20Time-to-event%20data_03112017_C26s2.pdf

Watkins, J. (2011, November). *Topic 15: Maximum likelihood estimation* [Lecture notes]. The University of Arizona. Retrieved July 17, 2025, from https://math.arizona.edu/~jwatkins/o-mle.pdf

Wilks, S. S. (1938). The large-sample distribution of the likelihood ratio for testing composite hypotheses. *The Annals of Mathematical Statistics*, *9*(1), 60–62.

Appendices

Appendix 1: GOIWBF World Record Data

i	Date (years)	R_i (s)	i	Date (years)	R_i (s)
1	2019.217	77.729	17	2021.264	65.841
2	2019.625	77.032	18	2021.353	65.254
3	2019.744	74.652	19	2021.375	64.922
4	2019.858	73.515	20	2021.436	62.922
5	2019.994	73.224	21	2021.856	62.893
6	2020.281	72.885	22	2022.061	62.703
7	2020.425	72.544	23	2022.072	62.623
8	2020.436	71.446	24	2022.153	61.896
9	2020.547	71.425	25	2022.572	61.356
10	2020.575	70.084	26	2022.600	61.205
11	2020.594	70.032	27	2022.908	61.061
12	2020.675	69.717	28	2022.961	60.962
13	2020.789	69.522	29	2023.194	60.907
14	2020.869	68.303	30	2023.217	60.152
15	2020.978	66.507	31	2023.253	59.885
16	2021.189	66.182	32	2025.450	58.326

Each record date t_i is obtained by subtracting the date of the first record, 2019.217, from the calendar date.

Appendix 2: Waiting Times Between Records

i	x_i	i	x_i	i	x_i	i	x_i
1	0.408	9	0.028	17	0.089	25	0.028
2	0.119	10	0.019	18	0.022	26	0.308
3	0.114	11	0.081	19	0.061	27	0.053
4	0.136	12	0.114	20	0.420	28	0.233
5	0.287	13	0.080	21	0.205	29	0.023
6	0.144	14	0.109	22	0.011	30	0.036
7	0.011	15	0.211	23	0.081	31	2.197
8	0.111	16	0.075	24	0.419		

Waiting times are in years. Since the waiting times telescope, $\sum_{i=1}^{31} x_i = t_{32} - t_1 = 6.233$.

Appendix 3: Exponential Model – Log-Likelihood and Second Derivative

$$\ell(\lambda) = \ln\Big(\lambda^{31} e^{-\lambda(x_1+x_2+\cdots+x_{31})}\Big) = \ln \lambda^{31} + \ln e^{-\lambda(x_1+x_2+\cdots+x_{31})} = 31 \ln \lambda - \lambda \sum_{i=1}^{31} x_i$$

$$\frac{d\ell}{d\lambda} = \frac{31}{\lambda} - \sum_{i=1}^{31} x_i = 0 \implies \lambda = \frac{31}{\sum_{i=1}^{31} x_i} = \frac{31}{6.233} \approx 4.974$$

$$\frac{d^2\ell}{d\lambda^2} = -\frac{31}{\lambda^2} = -\frac{31}{(4.974)^2} \approx -1.253 < 0$$

Appendix 4: Exponential Model – Expected Waiting Time

$$\begin{aligned}\mathbb{E}[X] &= \Big[-xe^{-4.974x} \Big]_0^\infty + \int_0^\infty e^{-4.974x}\, dx = \Big[-\frac{x}{e^{4.974x}} \Big]_0^\infty + \Big[-\frac{e^{-4.974x}}{4.974} \Big]_0^\infty \\ &= \lim_{x\to\infty} \left(-\frac{x}{e^{4.974x}} \right) - 0 + \lim_{x\to\infty} \left(-\frac{1}{4.974\, e^{4.974x}} \right) + \frac{1}{4.974} = \lim_{x\to\infty} \left(-\frac{x}{e^{4.974x}} \right) + \frac{1}{4.974}\end{aligned}$$

Appendix 5: Weibull Model – Log-Likelihood

$$\begin{aligned}\ell(\eta, \beta) &= \ln\left[\prod_{i=1}^n f(x_i \mid \eta, \beta) \right] = \sum_{i=1}^n \ln f(x_i \mid \eta, \beta) = \sum_{i=1}^n \left[\ln \frac{\beta}{\eta} + \ln\left(\frac{x_i}{\eta} \right)^{\beta-1} + \ln e^{-(x_i/\eta)^\beta} \right] \\ &= \sum_{i=1}^n \left[\ln \beta - \ln \eta + (\beta - 1)(\ln x_i - \ln \eta) - \left(\frac{x_i}{\eta} \right)^\beta \right] \\ &= n \ln \beta - n \ln \eta + \beta \sum_{i=1}^n \ln x_i - n\beta \ln \eta - \sum_{i=1}^n \ln x_i + n \ln \eta - \sum_{i=1}^n \left(\frac{x_i}{\eta} \right)^\beta \\ &= n \ln \beta - n\beta \ln \eta + (\beta - 1) \sum_{i=1}^n \ln x_i - \sum_{i=1}^n \left(\frac{x_i}{\eta} \right)^\beta\end{aligned}$$

Appendix 6: The Derivative of a Sum

For functions $f_1(x), \ldots, f_n(x)$ that are differentiable at x,

$$\begin{aligned}\frac{d}{dx} \left[\sum_{i=1}^n f_i(x) \right] &= \lim_{h\to 0} \frac{\sum_{i=1}^n f_i(x+h) - \sum_{i=1}^n f_i(x)}{h} = \lim_{h\to 0} \sum_{i=1}^n \frac{f_i(x+h) - f_i(x)}{h} \\ &= \sum_{i=1}^n \lim_{h\to 0} \frac{f_i(x+h) - f_i(x)}{h} = \sum_{i=1}^n \frac{d}{dx} \big[f_i(x) \big],\end{aligned}$$

since the limit of a finite sum is the sum of the limits.

Appendix 7: Weibull Model – Partial Derivative with Respect to β

Let $A_i = x_i/\eta$, which is constant with respect to β, so that $\frac{\partial}{\partial\beta}A_i^\beta = A_i^\beta \ln A_i$. Then

$$\begin{aligned}\frac{\partial \ell}{\partial \beta} &= \frac{n}{\beta} - n\ln\eta + \sum_{i=1}^{n}\ln x_i - \sum_{i=1}^{n}\left(\frac{x_i}{\eta}\right)^\beta \ln\frac{x_i}{\eta} \\ &= \frac{n}{\beta} - n\ln\eta + \sum_{i=1}^{n}\ln x_i - \frac{\sum_{i=1}^n x_i^\beta \ln x_i}{\eta^\beta} + \frac{\ln\eta \sum_{i=1}^n x_i^\beta}{\eta^\beta}.\end{aligned}$$

Substituting $\eta^\beta = \frac{1}{n}\sum_{i=1}^n x_i^\beta$, the last term equals $n\ln\eta$ and cancels the second term, giving

$$\frac{\partial \ell}{\partial \beta} = \frac{n}{\beta} + \sum_{i=1}^{n}\ln x_i - \frac{n\sum_{i=1}^n x_i^\beta \ln x_i}{\sum_{i=1}^n x_i^\beta}.$$

Appendix 8: Weibull Model – Expected Waiting Time

$$\mathbb{E}[X] = \int_0^\infty x\cdot\frac{\beta}{\eta}\left(\frac{x}{\eta}\right)^{\beta-1} e^{-(x/\eta)^\beta}\,dx = \frac{\beta}{\eta^\beta}\int_0^\infty x^\beta e^{-(x/\eta)^\beta}\,dx.$$

Let $u = (x/\eta)^\beta$, so that $x = \eta u^{1/\beta}$ and $dx = \frac{\eta}{\beta}u^{\frac{1}{\beta}-1}\,du$; the limits are unchanged. Then

$$\mathbb{E}[X] = \frac{\beta}{\eta^\beta}\int_0^\infty \eta^\beta u\, e^{-u}\cdot\frac{\eta}{\beta}u^{\frac{1}{\beta}-1}\,du = \eta\int_0^\infty u^{1/\beta}e^{-u}\,du = \eta\,\Gamma\!\left(1+\frac{1}{\beta}\right).$$

Appendix 9: Cumulative Distribution Functions and Interval Bounds

Exponential.

$$F_X(x) = 0.95 \implies e^{-4.974x} = 0.05 \implies x = -\frac{\ln 0.05}{4.974} \approx 0.602$$

Weibull. Using the substitution $u = (s/\eta)^\beta$, $du = \frac{\beta}{\eta}\left(\frac{s}{\eta}\right)^{\beta-1}ds$,

$$F(x) = \int_0^x \frac{\beta}{\eta}\left(\frac{s}{\eta}\right)^{\beta-1} e^{-(s/\eta)^\beta}\,ds = \int_0^{(x/\eta)^\beta} e^{-u}\,du = 1 - e^{-(x/\eta)^\beta},$$

$$F(x) = 0.95 \implies \left(\frac{x}{\eta}\right)^\beta = -\ln 0.05 \implies x = \eta(-\ln 0.05)^{1/\beta} \approx 0.172\times 3.899 \approx 0.670.$$