Risk and Uncertainty

BEE 6940 Lecture 2

January 30, 2023

Any Questions?

Table of Contents

What Is Climate Risk?

What Is Climate Risk?

Climate risk: "risk" created or enhanced by the impacts of climate change

What Is Climate Risk?

Climate risk: "risk" created or enhanced by the impacts of climate change

Strong interactions between these impacts and broader socioeconomic dynamics results in complex dynamics.

Climate Impacts are Diverse

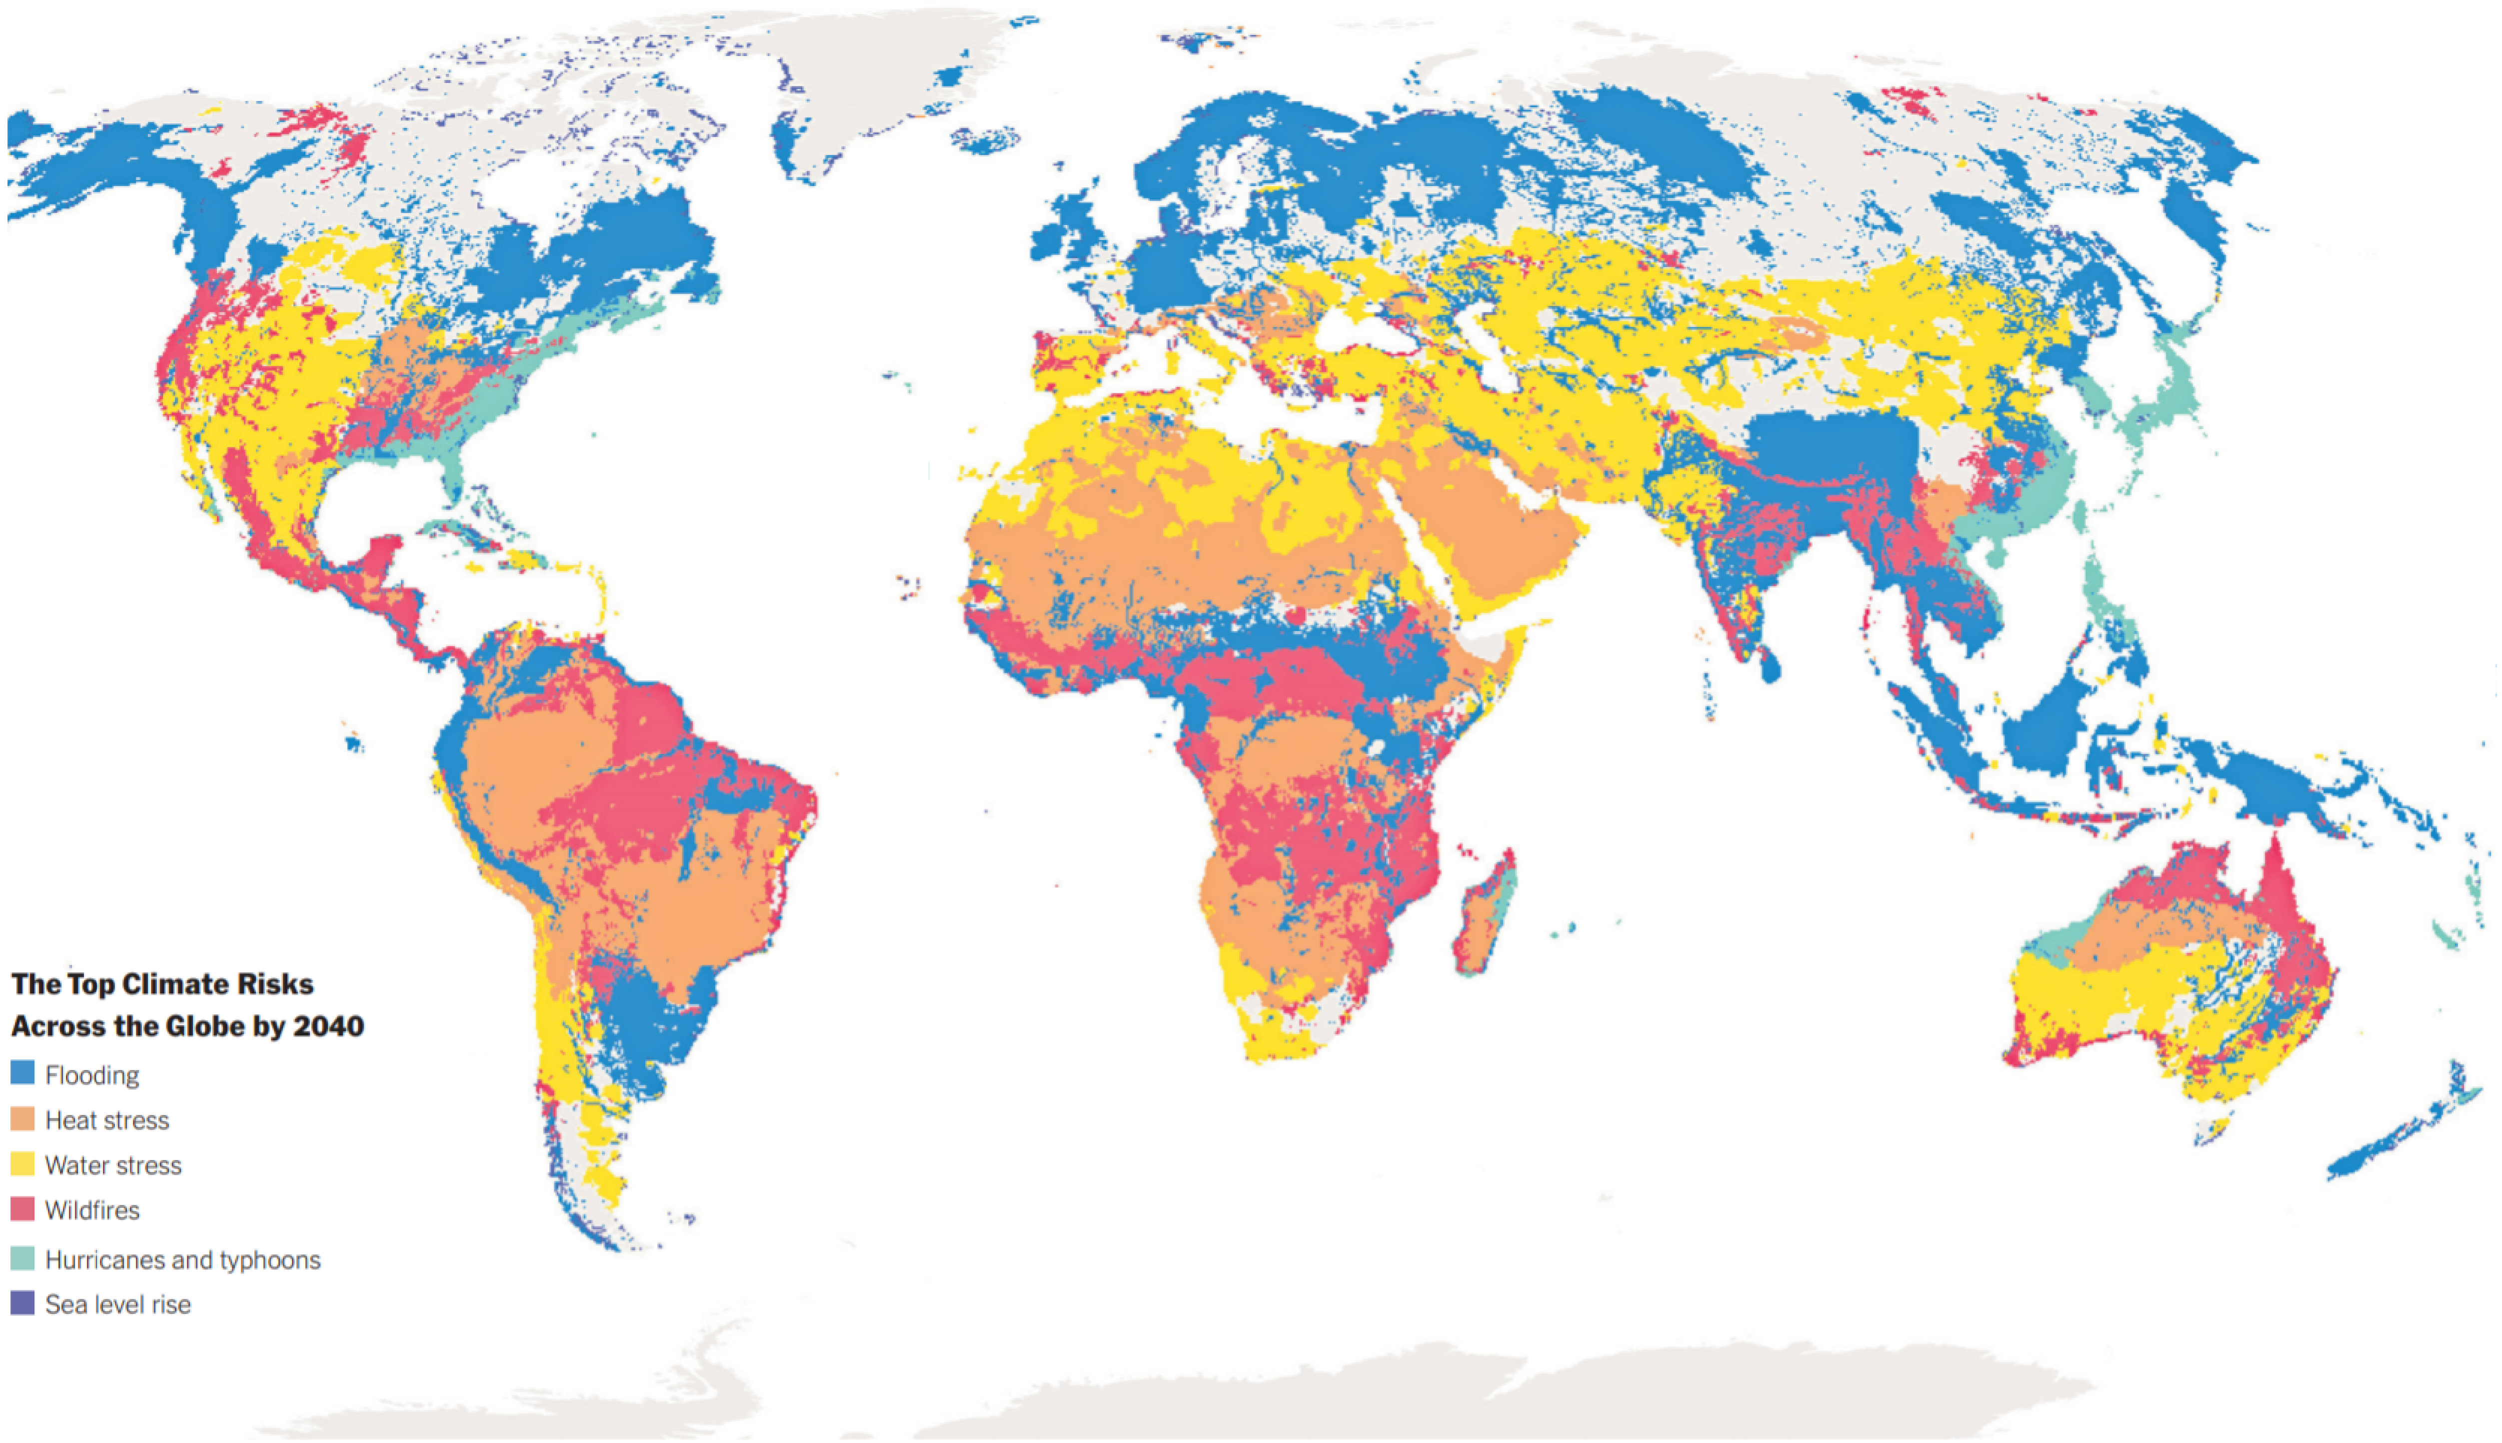

Source: Four Twenty Seven and the New York Times

Climate Risks are Worsening

Climate crisis

Extreme Weather | Climate | Weather | Capital Weather Gang | Environment | Climate Lab

WEATHER

How climate change will make atmospheric rivers even worse

Atmospheric rivers are projected to become wetter, larger and more damaging as temperatures rise

SUNDAY MORNING ›

Coastal residents on climate change: "The ocean's coming for you"

Relentless rain, record heat: study finds climate crisis worsened extreme weather

Scientists describe as 'very alarming' research that shows severe weather events were made more likely by climate change

NEWS | Climate Change | Add Topic +

Another above-average wildfire season for 2022. How climate change is making fires harder to predict and fight.

What Is Risk?

What Is Risk?

Intuitively: "Risk" is the possibility of loss, damages, or harm.

$$\text{Risk} = \text{Probability of Hazard} \times \text{Damages From Event}$$

What Is Risk?

Intuitively: "Risk" is the possibility of loss, damages, or harm.

$$\text{Risk} = \text{Probability of Hazard} \times \text{Damages From Event}$$

Things we don't think of as "risk":

- Good or neutral outcomes
- Deterministic outcomes

Some Cartoons About Risk

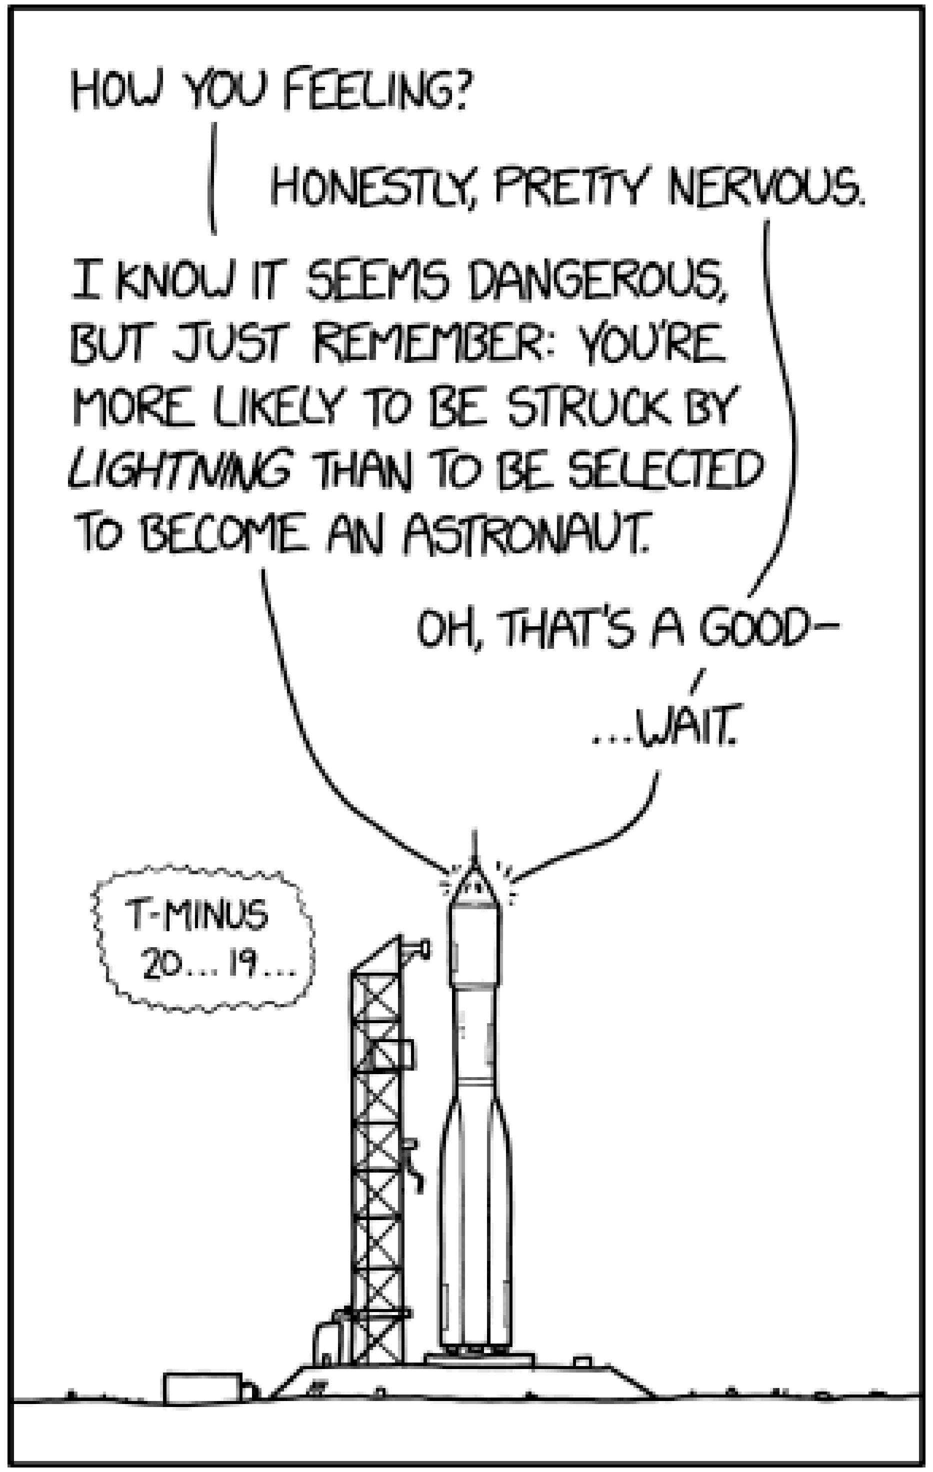

Source: *XKCD 2107*

Some Cartoons About Risk

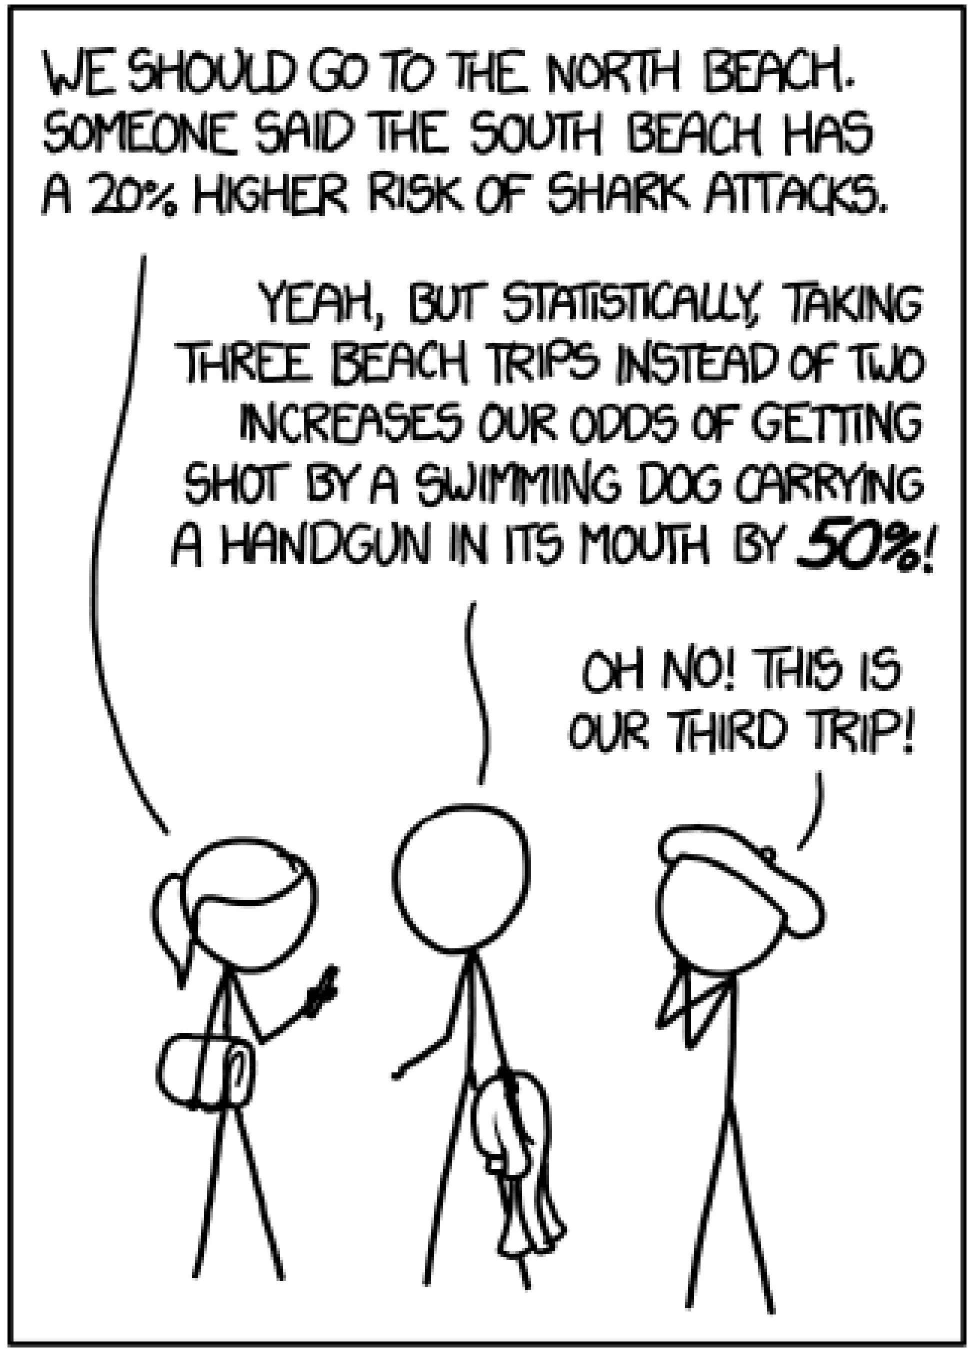

REMINDER: A 50% INCREASE IN A TINY RISK IS *STILL TINY.*

Source: XKCD 1252

What Is Risk?

Common framework:

Risk as a combination of

- Hazard
- Exposure
- Vulnerability
- *Response* (Simpson et al (2021))

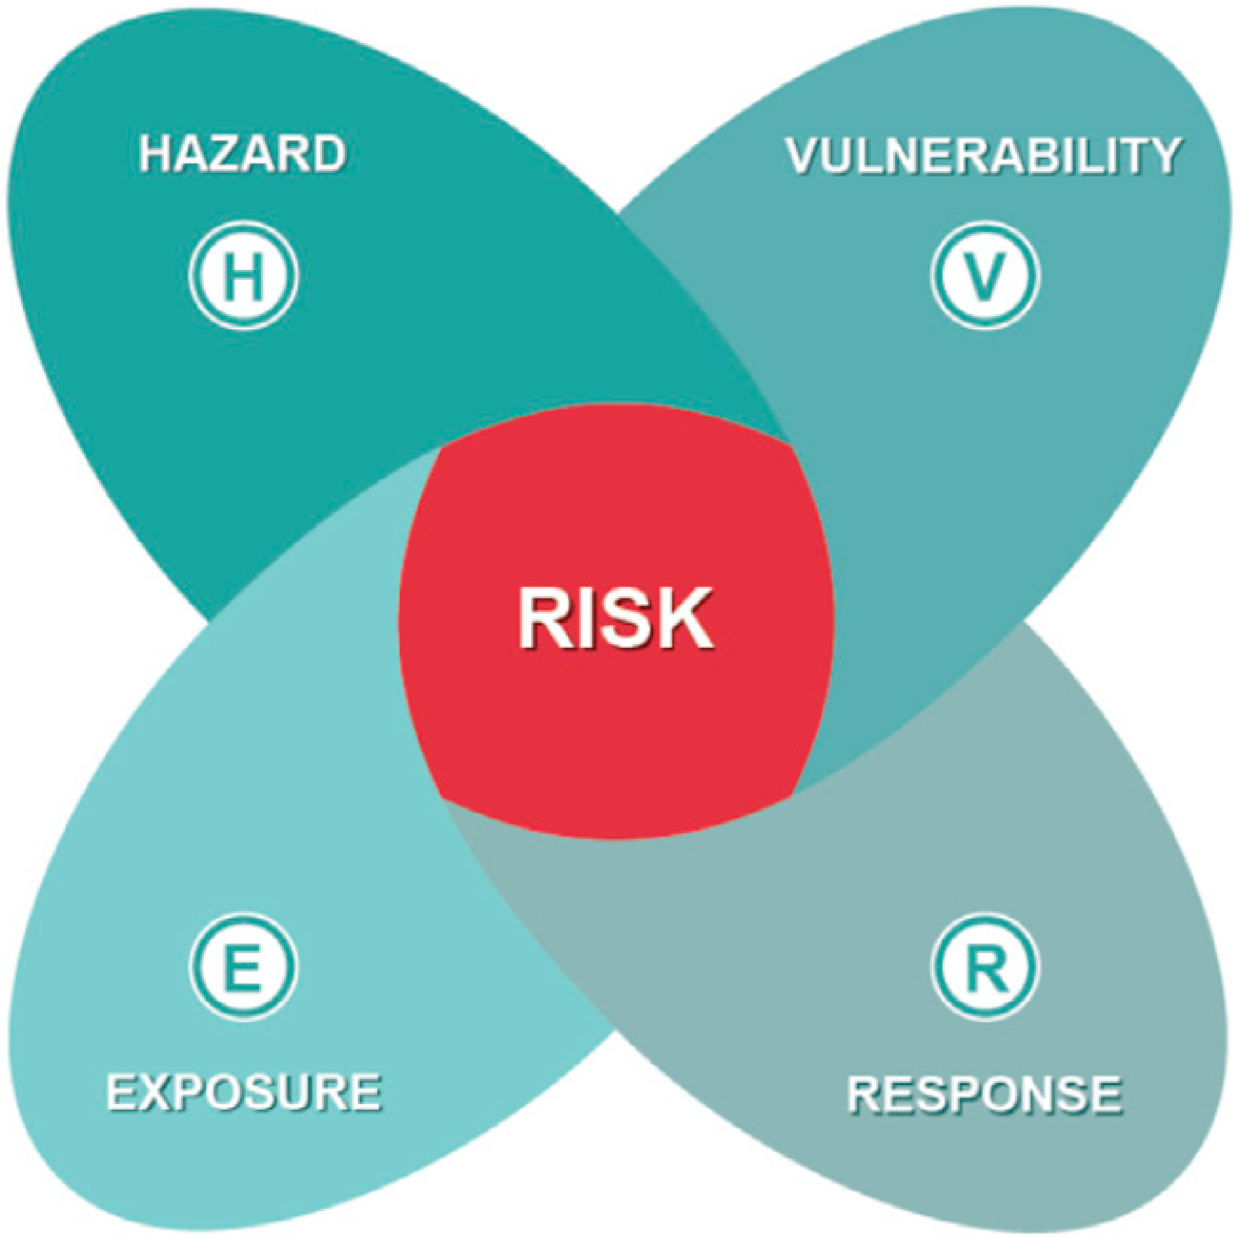

Source: *Simpson et al (2021)*

Defining Climate Risk

Climate Risk: Changes in risk stemming from the impacts of or response to climate change.

Defining Climate Risk

Climate Risk: Changes in risk stemming from the impacts of or response to climate change.

Hazards

- Drought/flooding
- Extreme temperatures
- Sea level rise
- Others!

Exposure/Vulnerability

- Compound events
- Urbanization
- Land Use, Land Cover Change

Motivating Questions

1. What are the potential impacts of climate change?
2. What can we say about their uncertainties?
3. What are the impacts of those uncertainties on the performance of risk-management strategies?

Uncertainty and Probability

Uncertainty and Risk Analysis

Uncertainty enters into the hazard-exposure-vulnerability-response model in a few ways:

- Uncertain hazards
- Uncertainty in model estimates of exposure or vulnerability
- Uncertainty in responses

But...

What exactly do we mean by uncertainty?

Glib answer: Uncertainty is a lack of certainty!

But...

What exactly do we mean by uncertainty?

Glib answer: Uncertainty is a lack of certainty!

Maybe better: Uncertainty refers to an inability to exactly describe current or future states.

Two Categories of Uncertainty

- **Aleatory Uncertainty**: Uncertainty resulting from *inherent randomness*
- **Epistemic Uncertainty**: Uncertainty resulting from *lack of knowledge*

Two Categories of Uncertainty

- **Aleatory Uncertainty**: Uncertainty resulting from *inherent randomness*
- **Epistemic Uncertainty**: Uncertainty resulting from *lack of knowledge*

The lines between aleatory and epistemic uncertainty are not always clear! This has implications for modeling and risk analysis.

On Epistemic Uncertainty

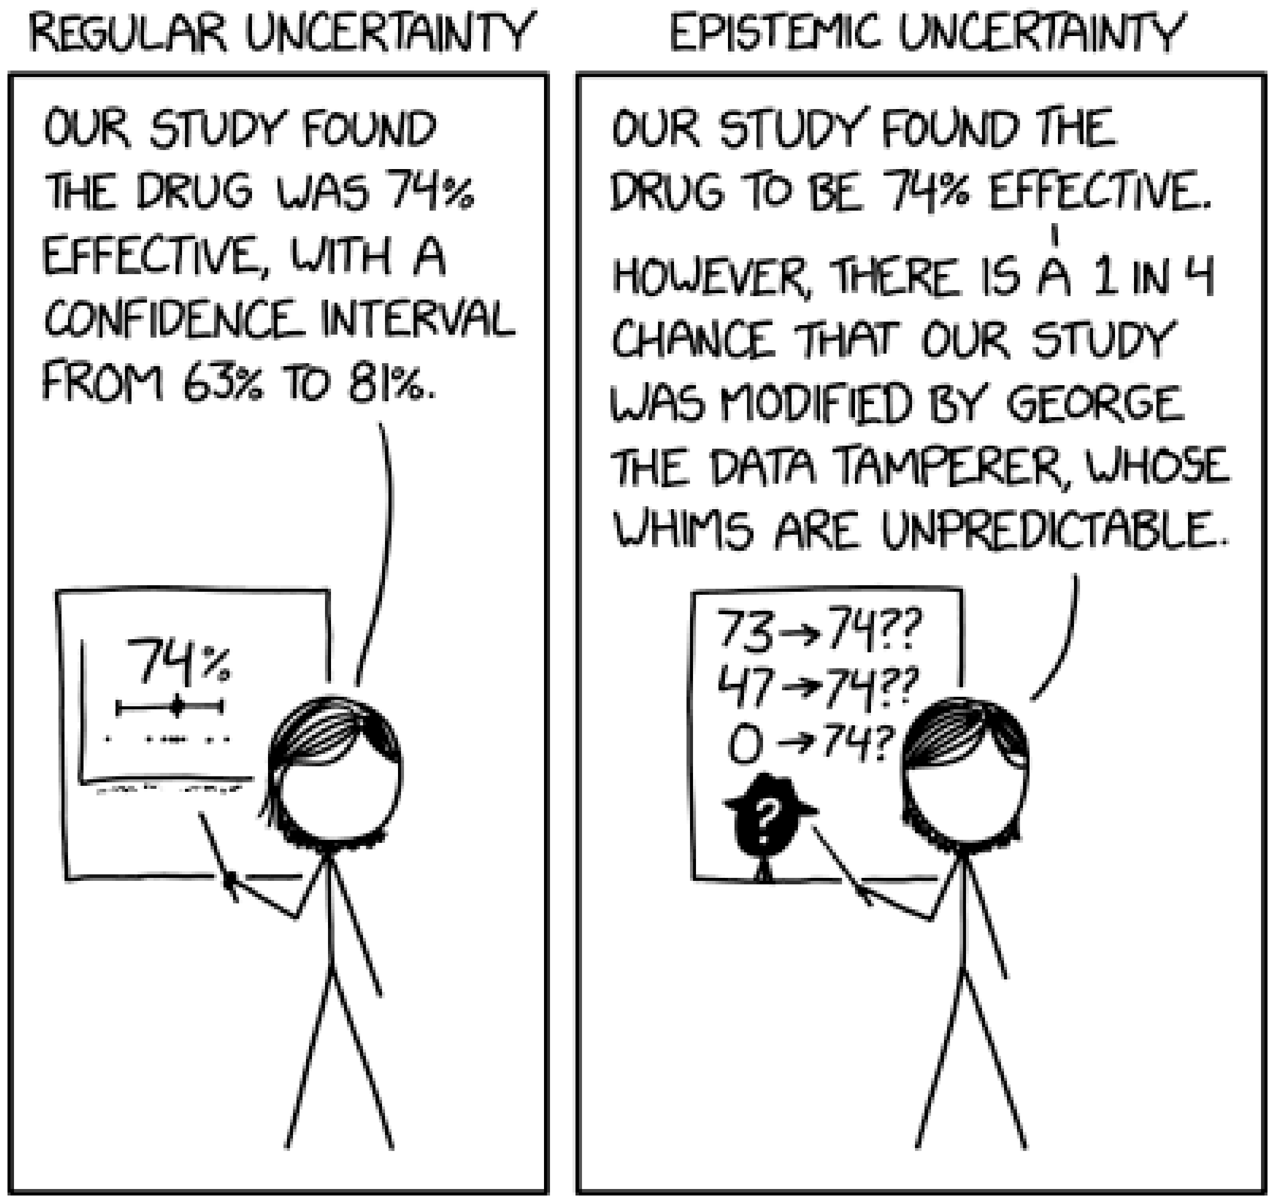

Source: XKCD 2440

Uncertainty and Probability

We often represent or describe uncertainties in terms of *probabilities*:

- Long-run frequency of an event (**frequentist**)
- Degree of belief that a proposition is true (**Bayesian**)

Confidence vs. Credible Intervals

The difference between the frequentist and Bayesian perspectives can be illustrated through the difference in how both conceptualize uncertainty in estimates.

Confidence vs. Credible Intervals

A Bayesian **credible interval** for some random quantity is conceptually straightforward:

An α-credible interval is an interval with an α% probability of containing the realized or "true" value.

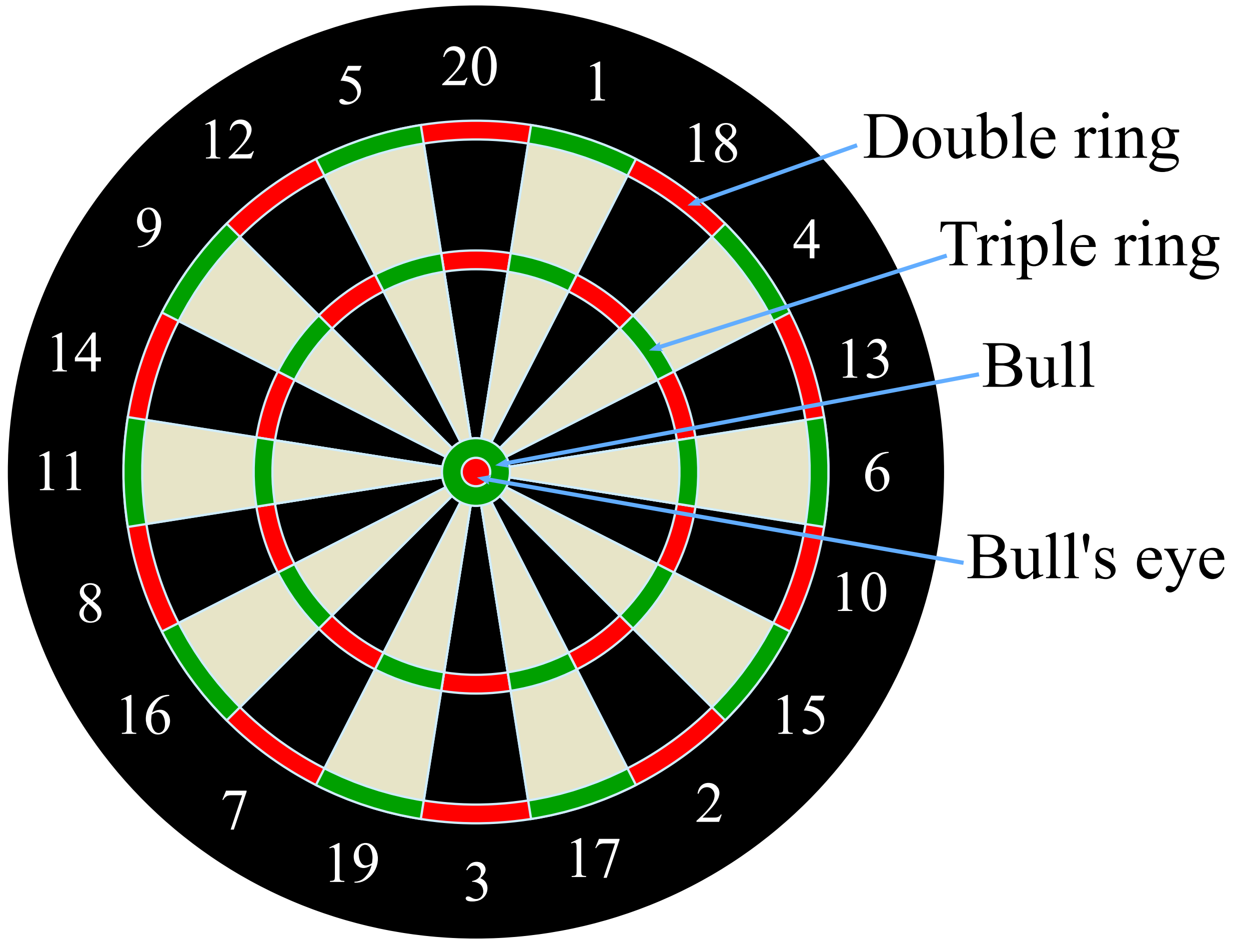

Source: Wikipedia

Confidence vs. Credible Intervals

However, this notion breaks down with the frequentist viewpoint: there is some "true value" for the associated estimate based on long-run frequencies.

With this view, it is incoherent to talk about probabilities corresponding to parameters. Instead, the key question is how frequently (based on repeated analyses of different datasets) your estimates are "correct".

Confidence vs. Credible Intervals

In other words, the confidence level $\alpha\%$ expresses the *pre-experimental* frequency by which a confidence interval will contain the true value.

So for a 95% confidence interval, there is a 5% chance that a given sample was an outlier and the interval is inaccurate.

Confidence vs. Credible Intervals

To understand frequentist **confidence intervals**, think of horseshoes! The post is a fixed target, and my accuracy as a horseshoe thrower captures how confident I am that I will hit the target with any given toss.

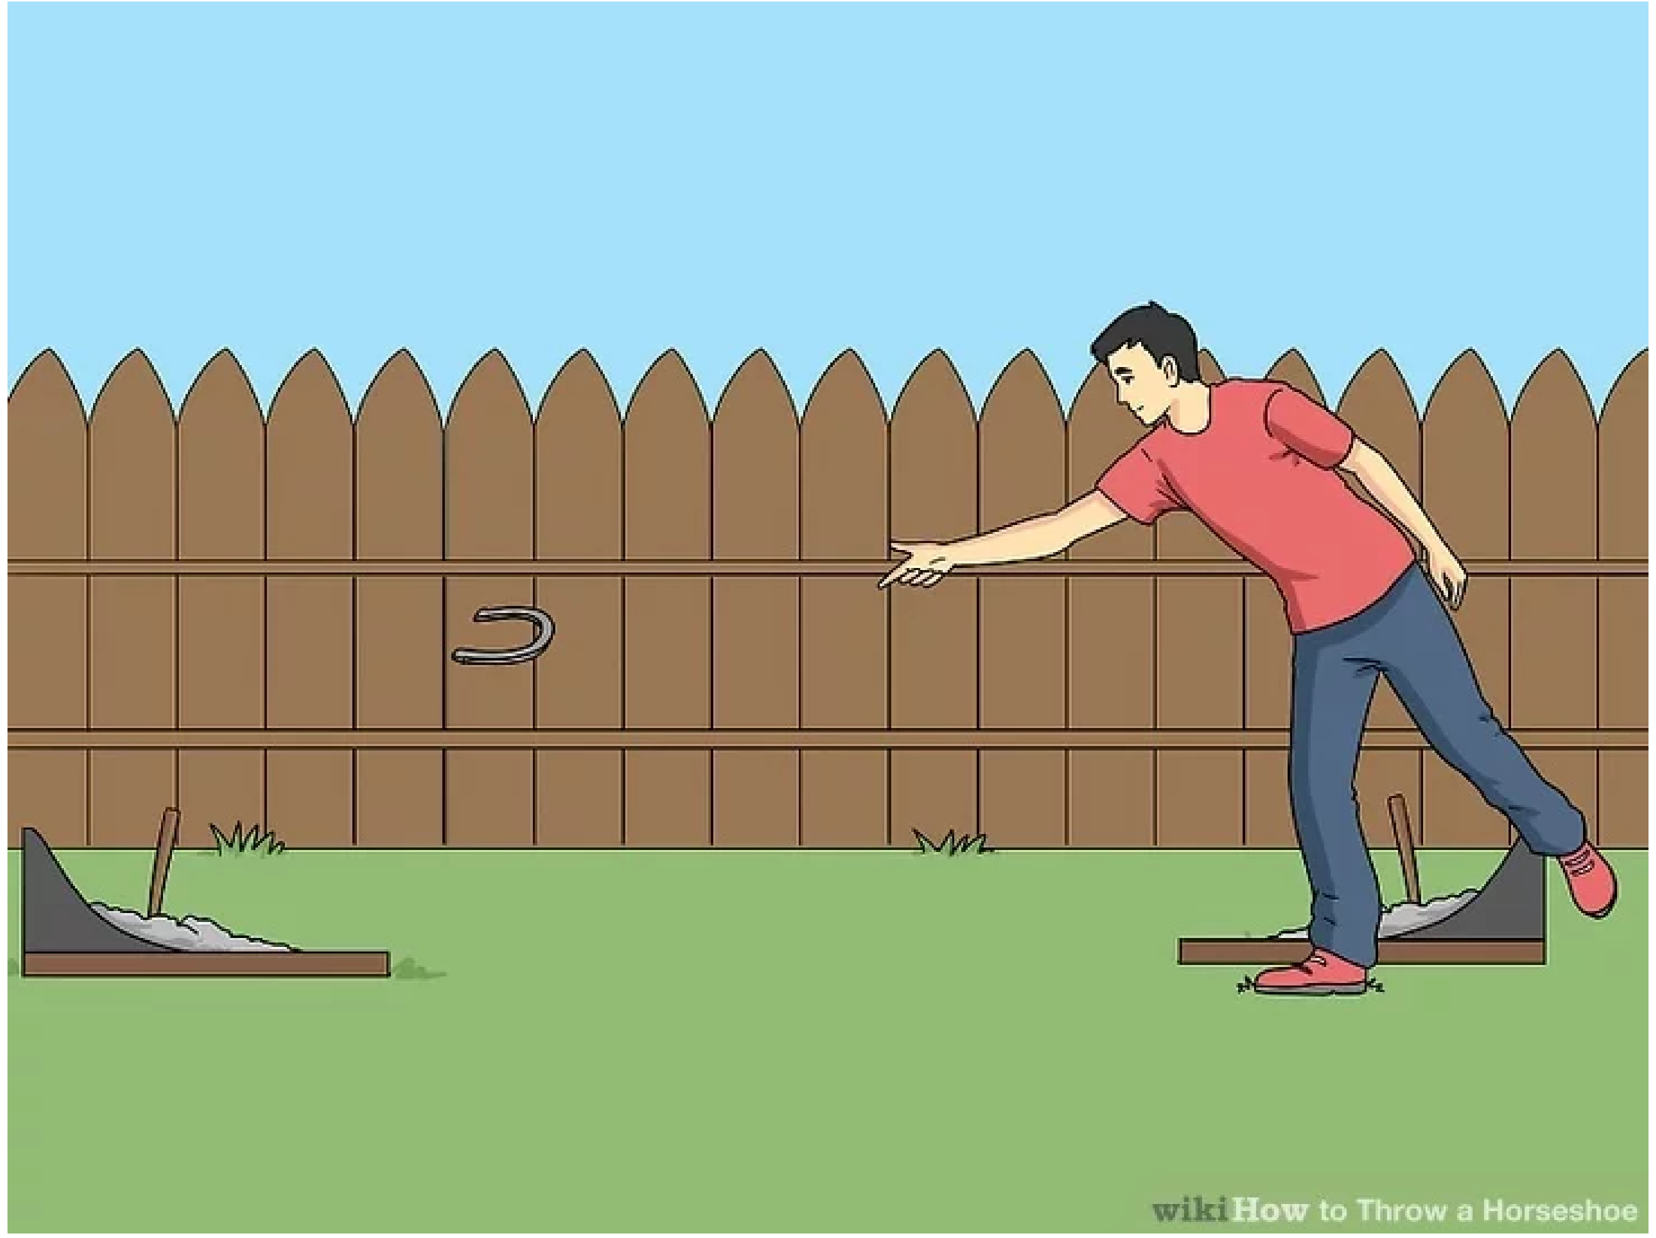

Source: https://www.wikihow.com/Throw-a-Horseshoe

Confidence vs. Credible Intervals

But once I make the throw, I've either hit or missed.

Generating a confidence interval is like throwing a horseshoe with a certain (pre-experimental) degree of accuracy.

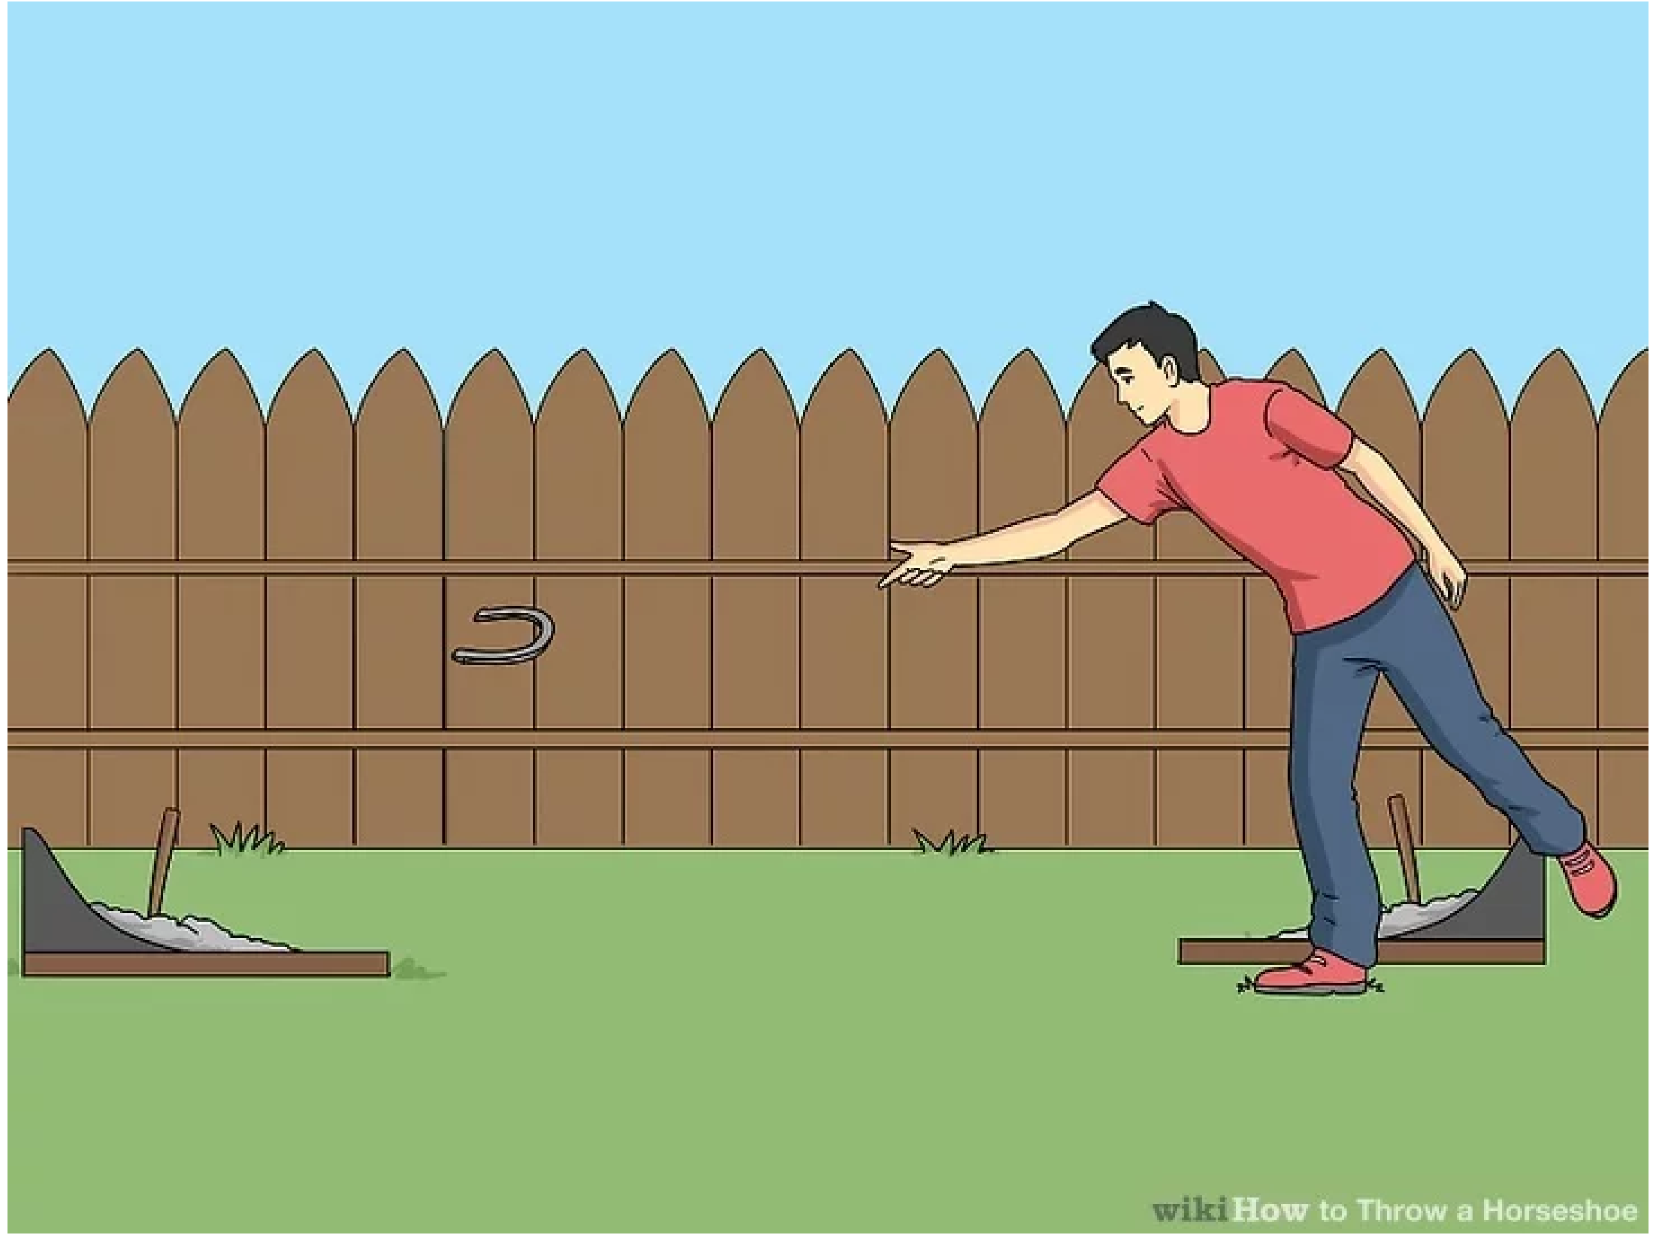

Source: https://www.wikihow.com/Throw-a-Horseshoe

Probability Distributions

Probabilities are often represented using a probability distribution, which are parameterized by a *probability density function*.

- Normal (Gaussian) Distribution: mean μ, variance σ^2
- Poisson Distribution: rate λ
- Binomial Distribution: # trials n, probability of success p
- Generalized Extreme Value Distribution: location μ, scale σ, shape ξ

Probability Models

A key consideration in uncertainty and risk analysis is defining an appropriate *probability model* for the data.

Many "default" approaches, such as linear regression, assume normal distributions and independent and identically-distributed residuals.

Deviations from Normality

Some typical ways in which these assumptions can fail:

- **skew** (more samples on one side of the mean than the other)

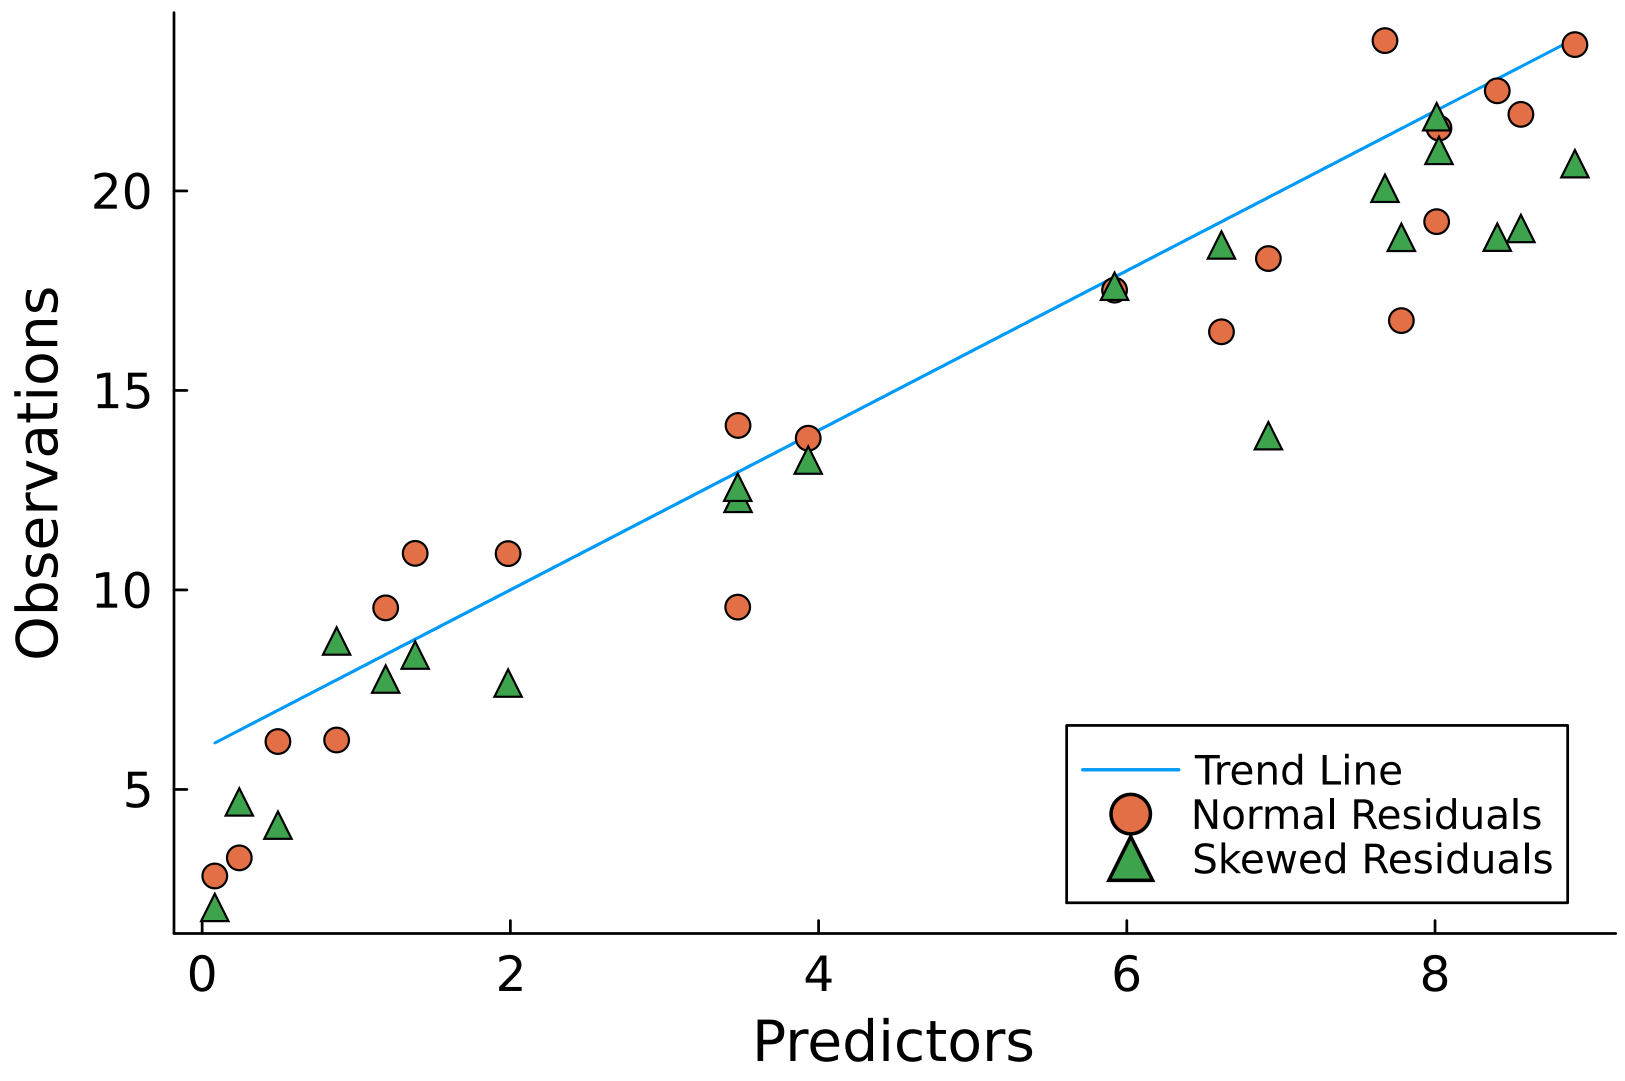

Deviations from Normality

Some typical ways in which these assumptions can fail:

- skew
- **fat tails** (probability of extremes)

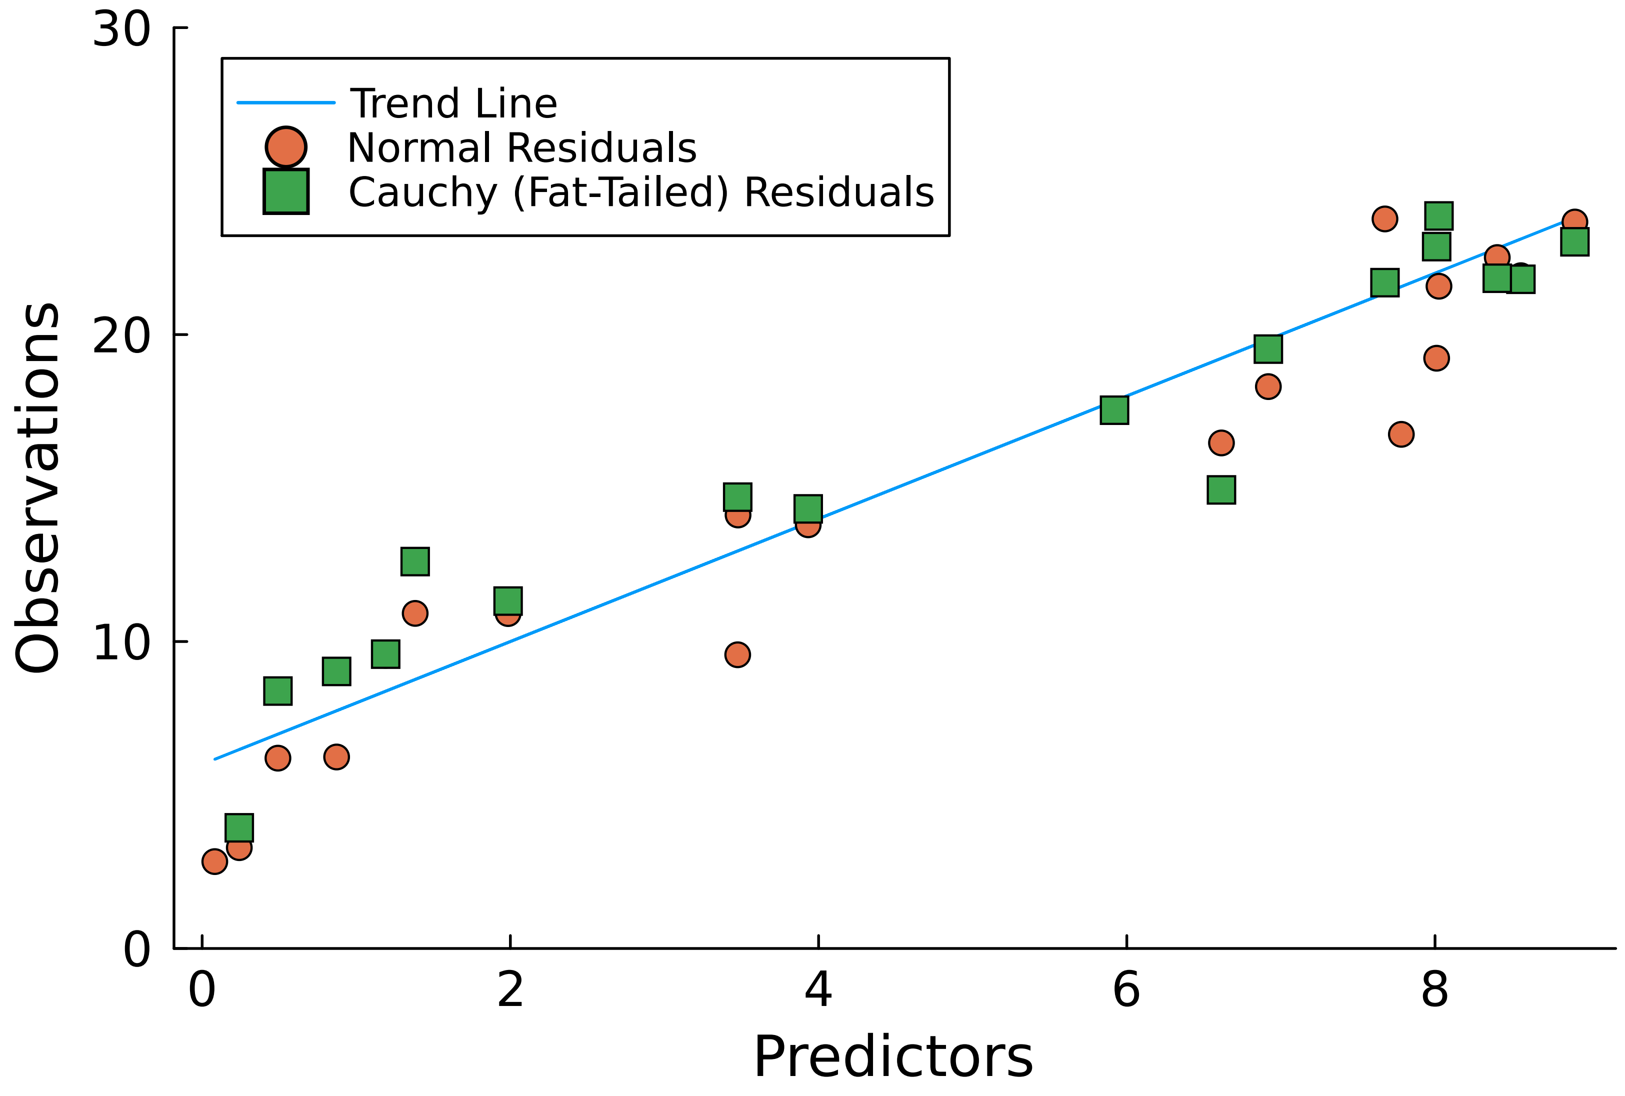

Deviations from Normality

Some typical ways in which these assumptions can fail:

- skew
- fat tails
- **(auto-)correlations**

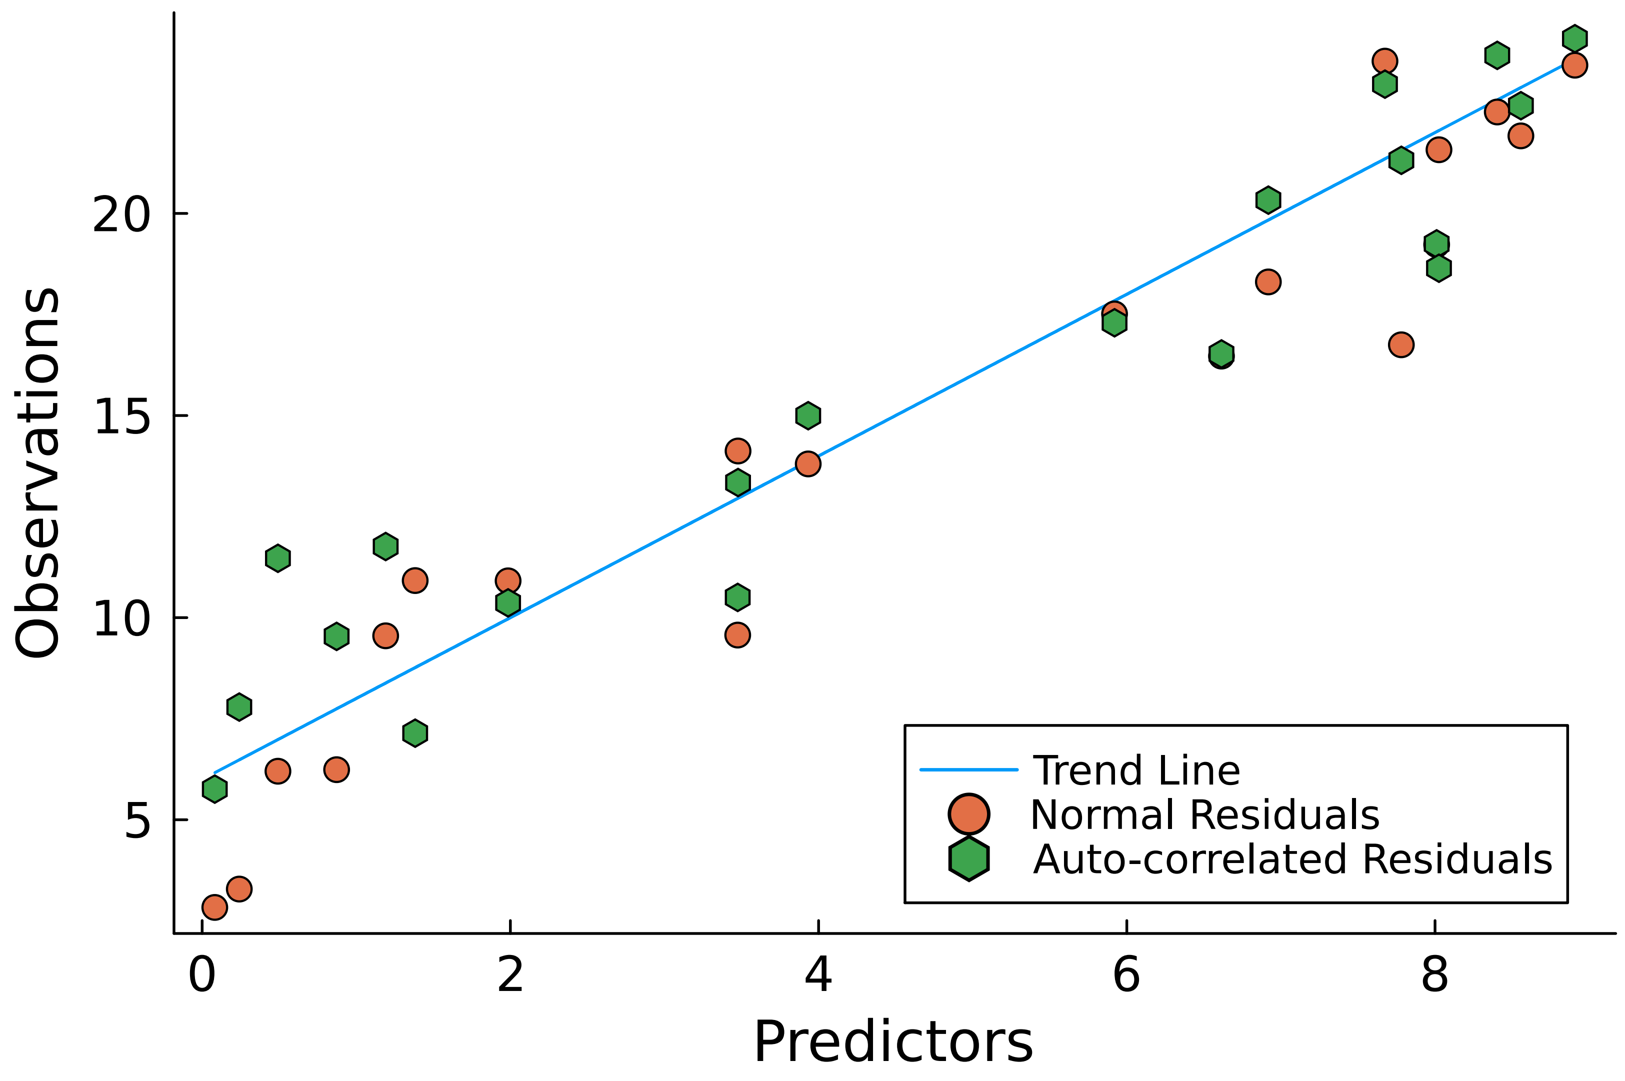

Diagnosing Quality of Fit

How can we know if a proposed probability model is appropriate for a data set?

Diagnosing Quality of Fit

Visual inspection often breaks down: our brains are very good at imposing structure (look up "gestalt principles").

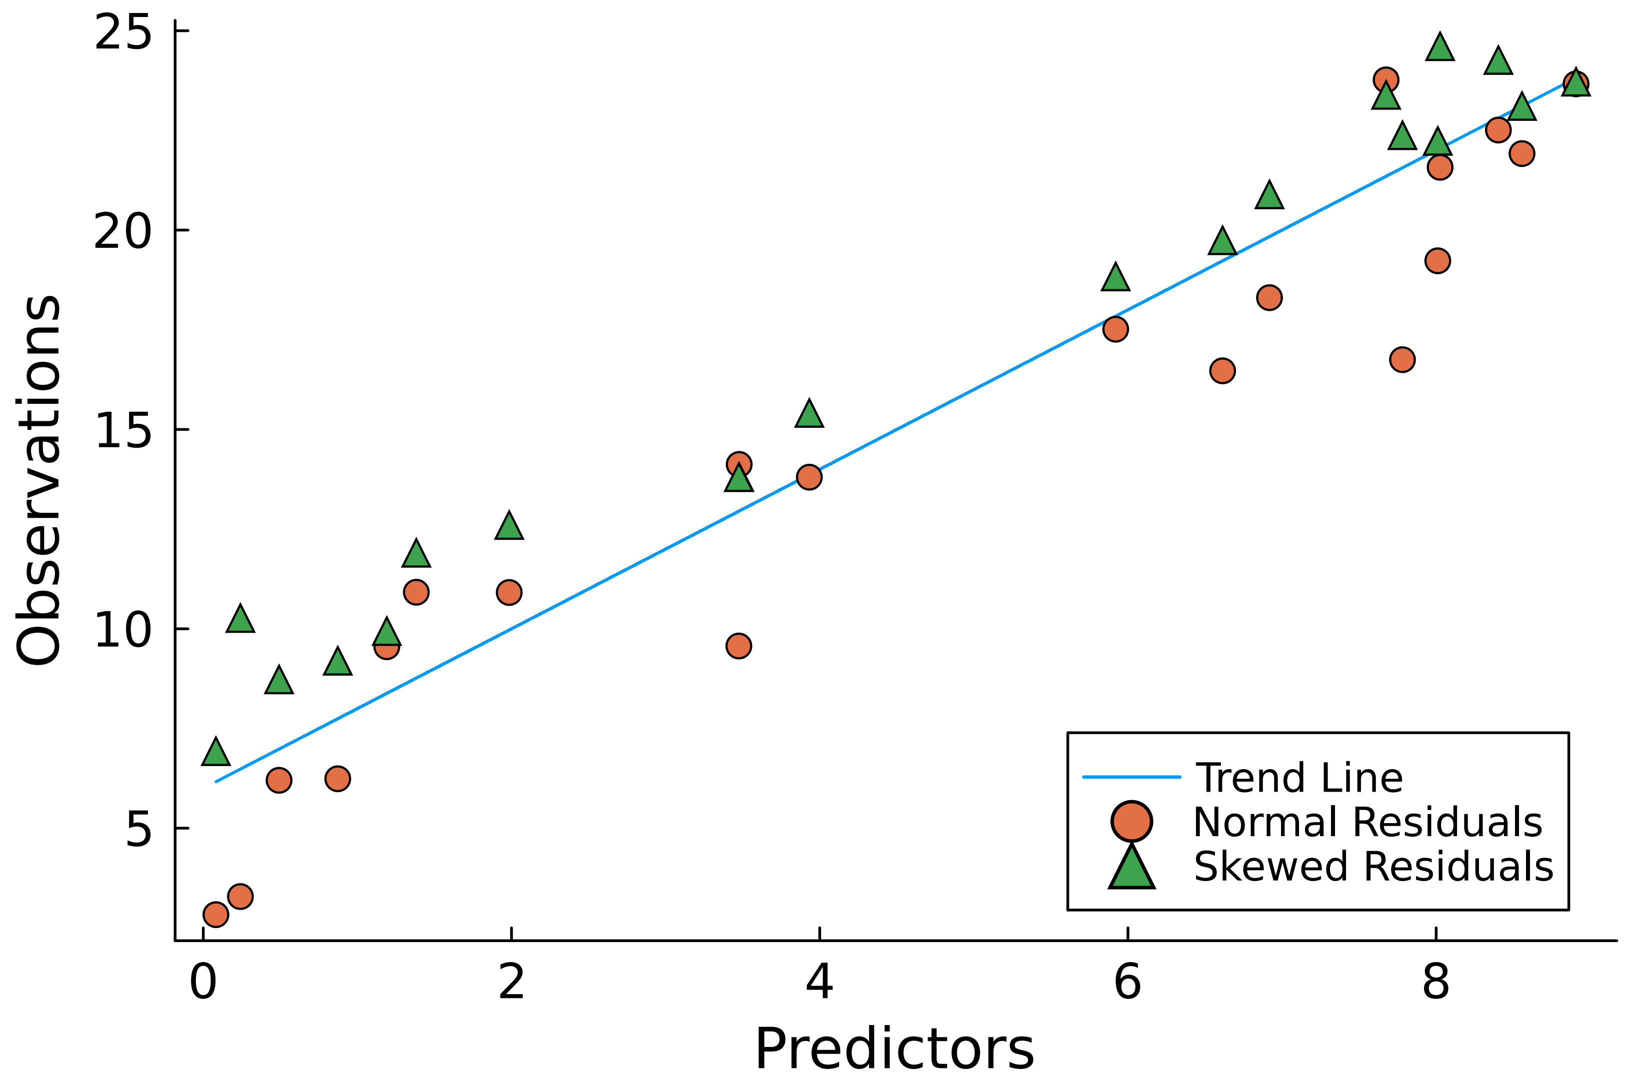

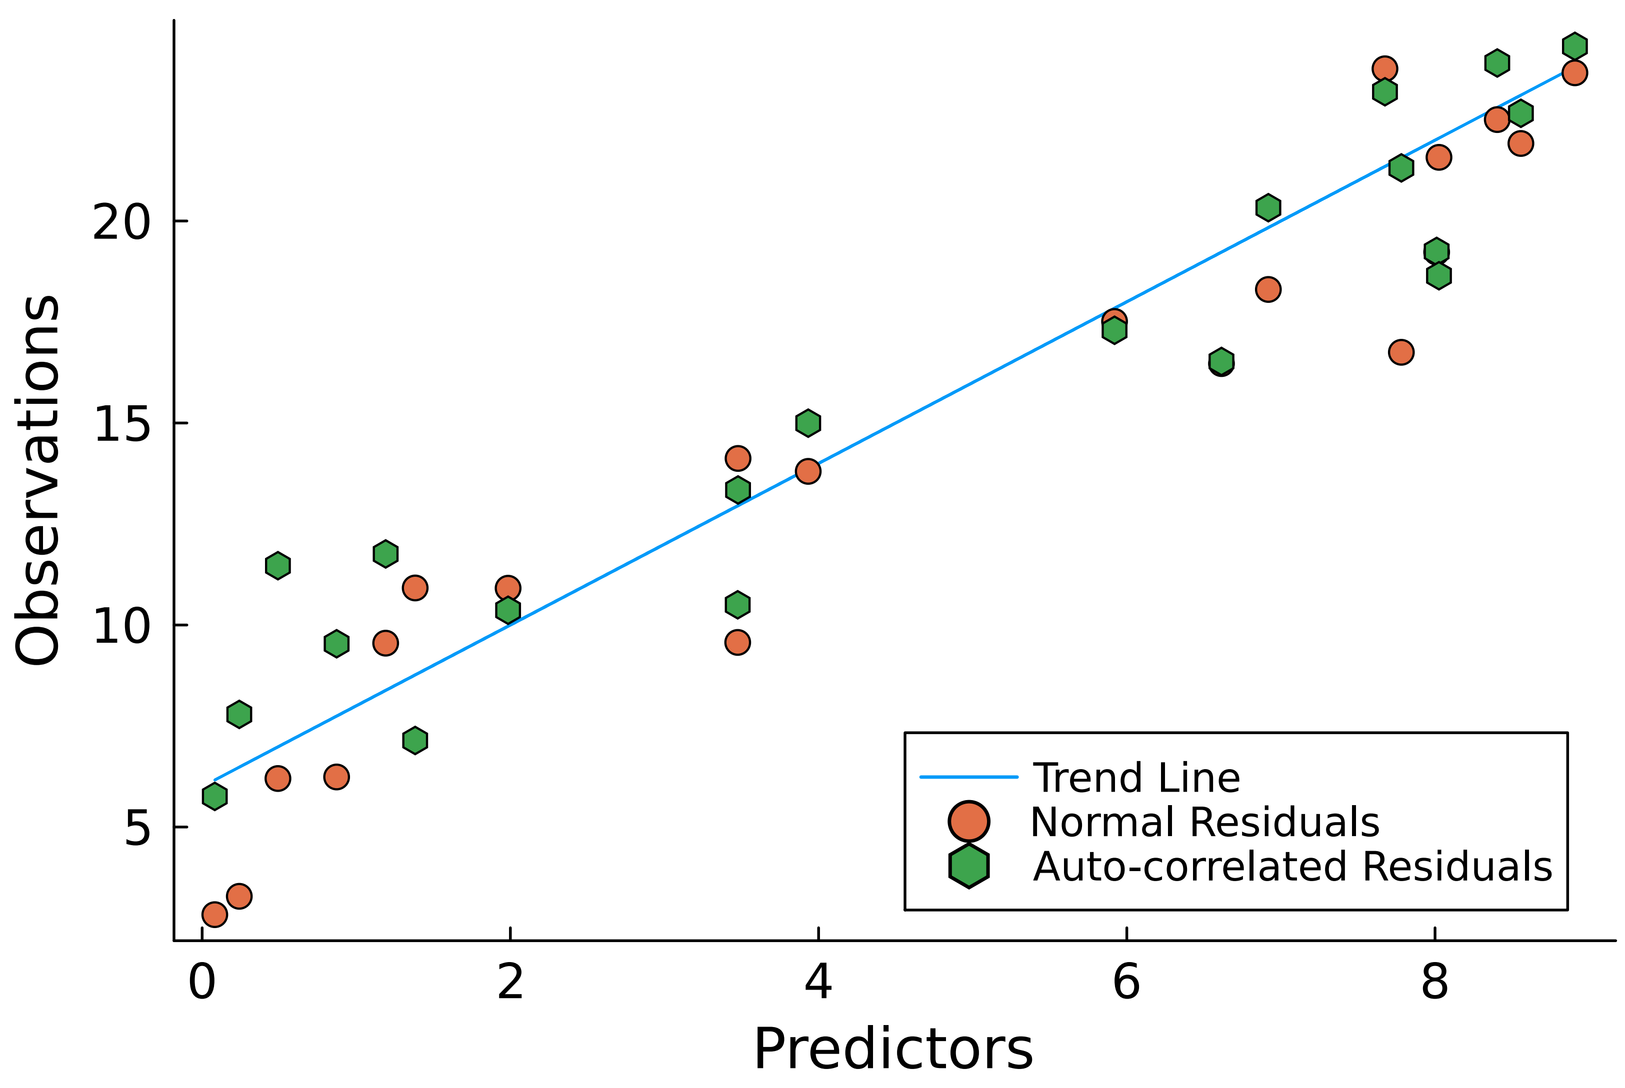

Quantile-Quantile Plots

One useful tool is a **quantile-quantile (Q-Q) plot**, which compares quantiles of two distributions.

If the quantiles match, the points will be roughly along the diagonal line, *e.g.* this comparison of normally-distributed data with a normal distribution.

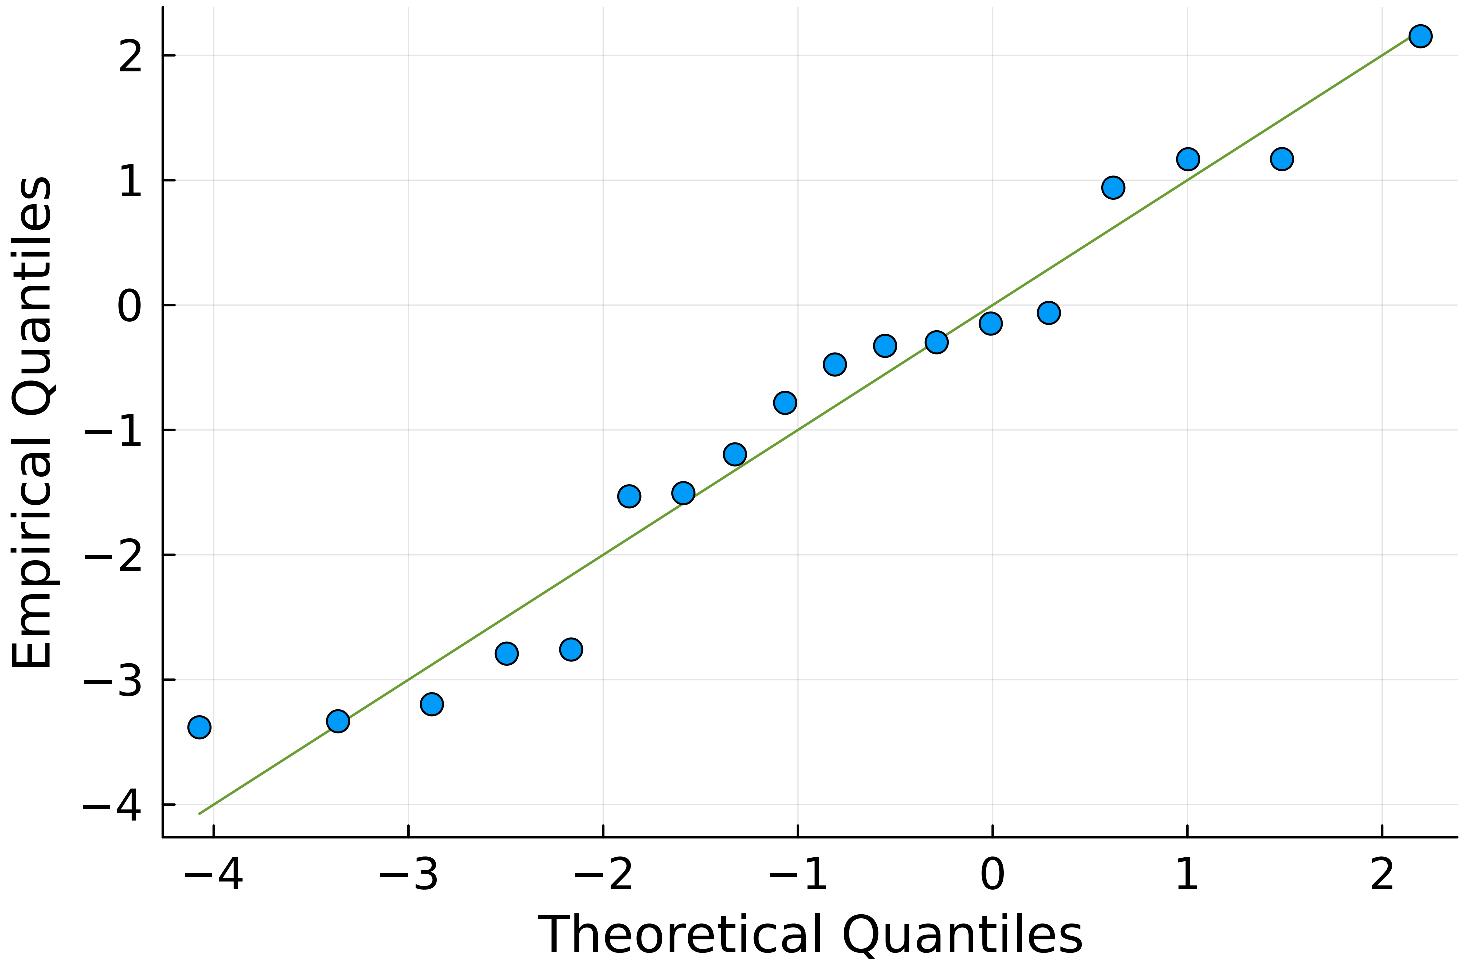

Quantile-Quantile Plots

If the points are below/above the 1:1 line, the theoretical distribution is over/under-predicting the associated quantiles.

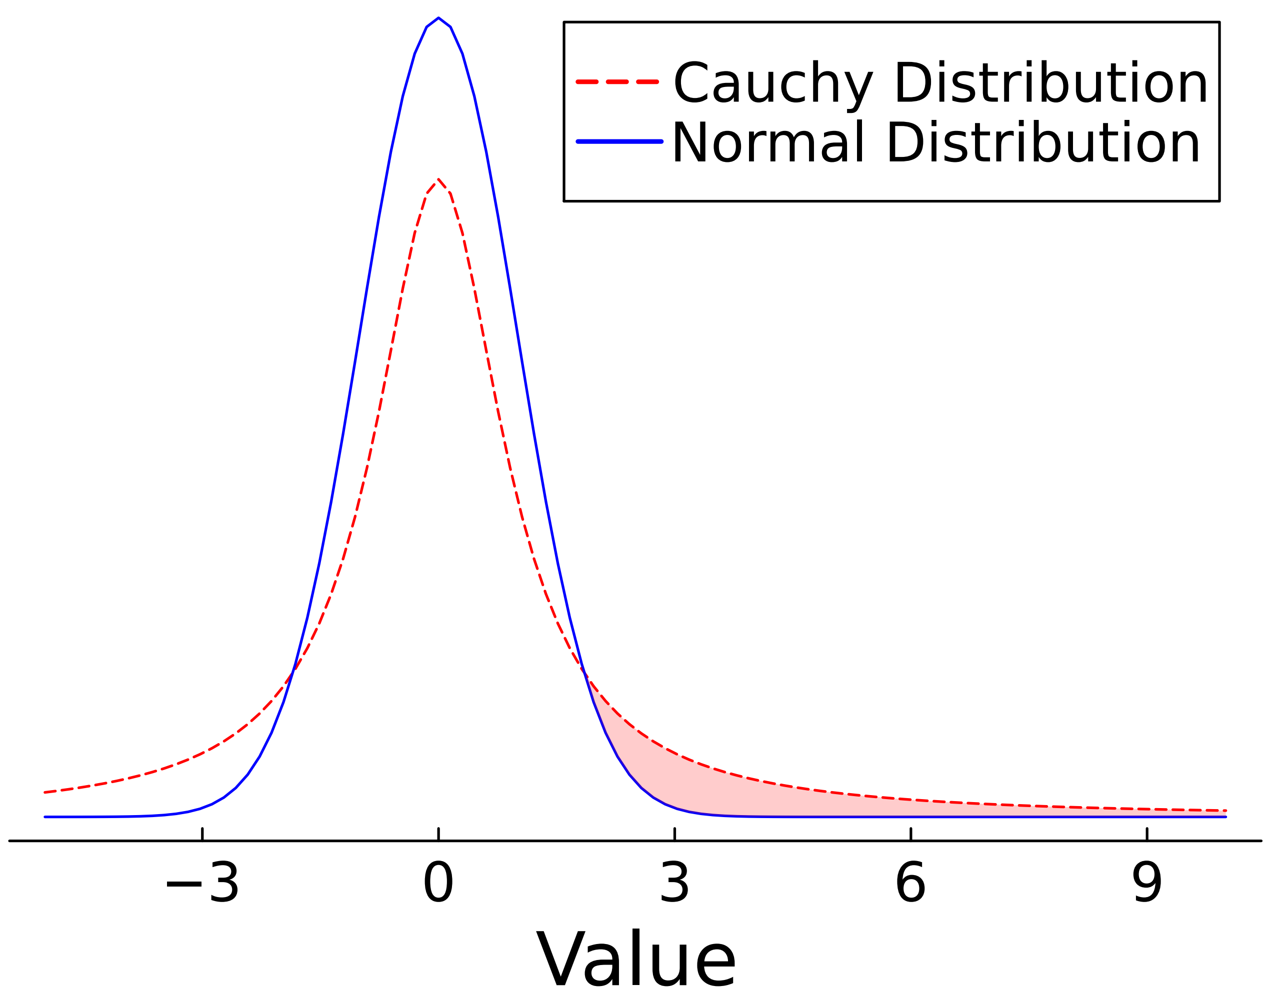

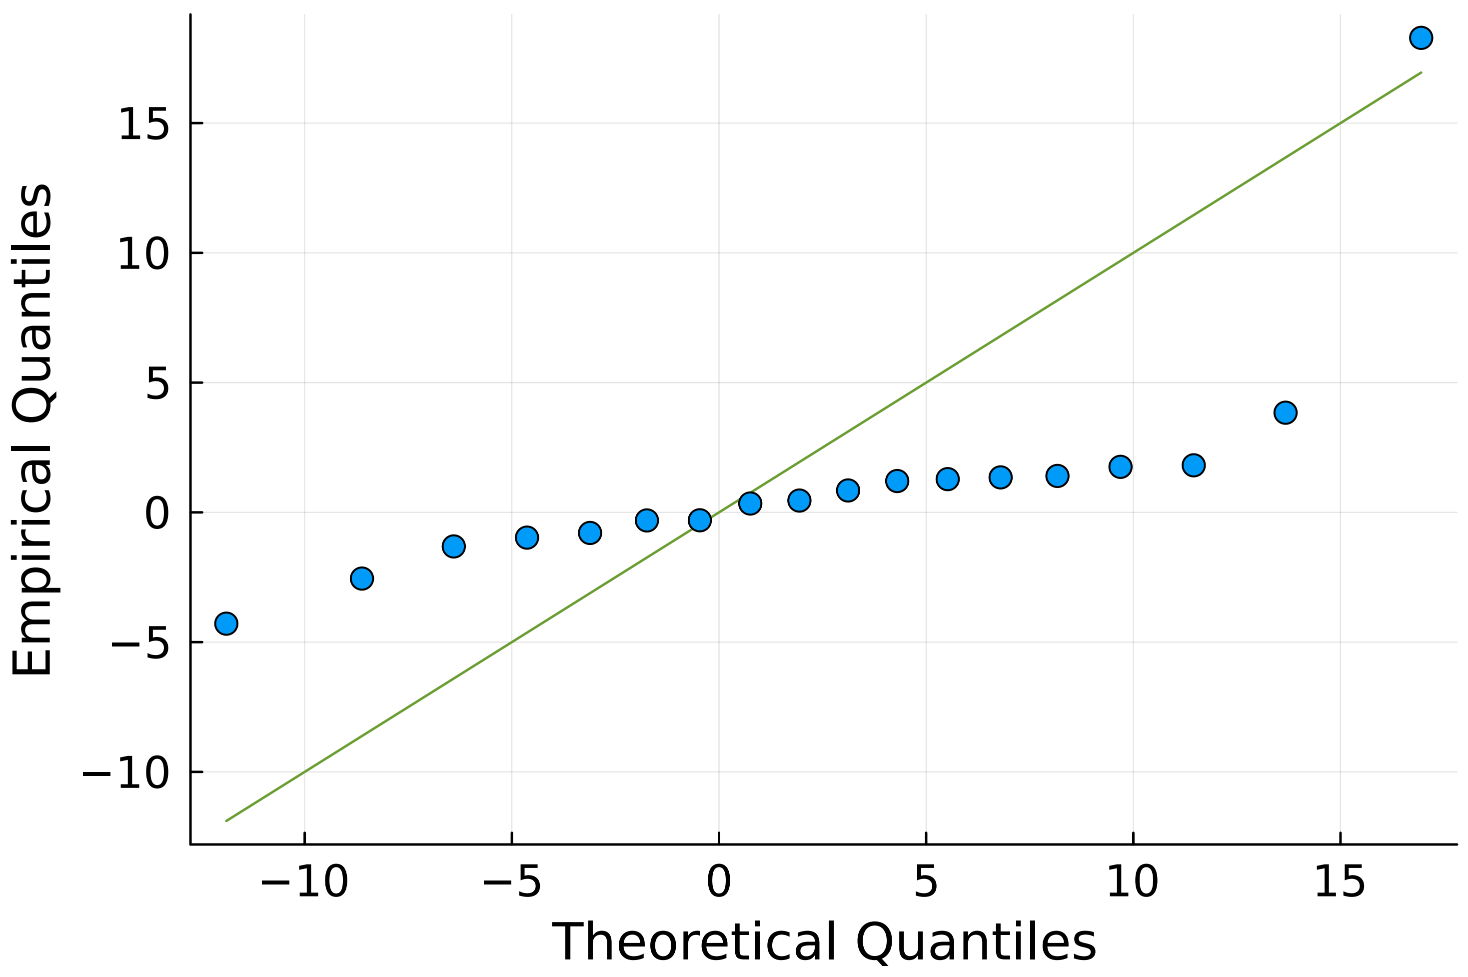

Cumulative Distribution Functions

Q-Q plots show similar information to a **Cumulative Distribution Function (CDF) plot**.

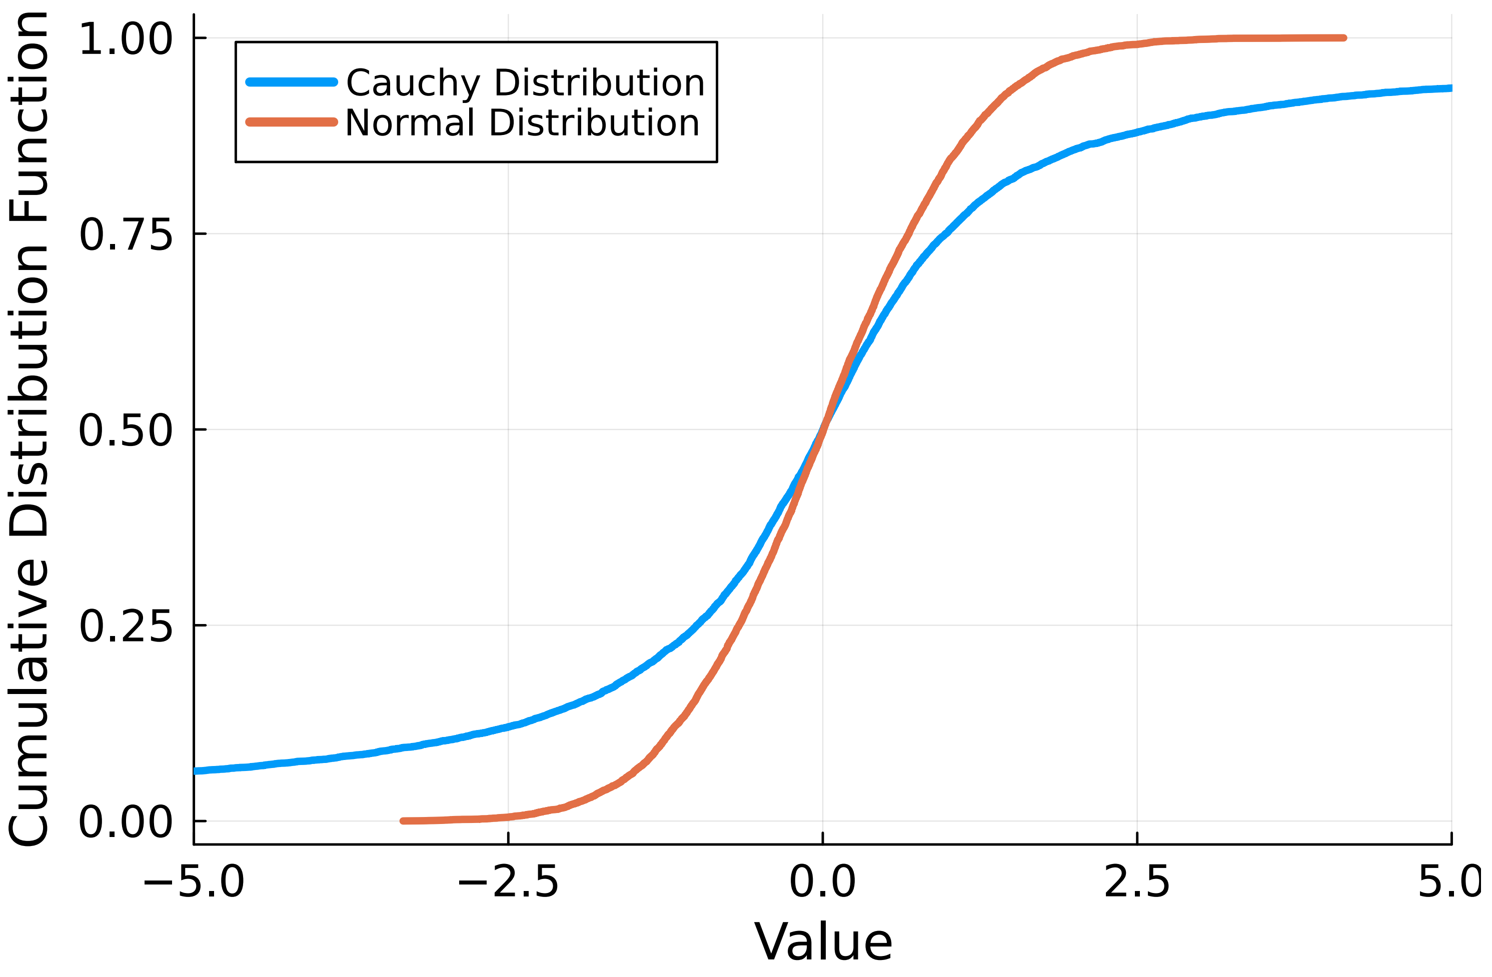

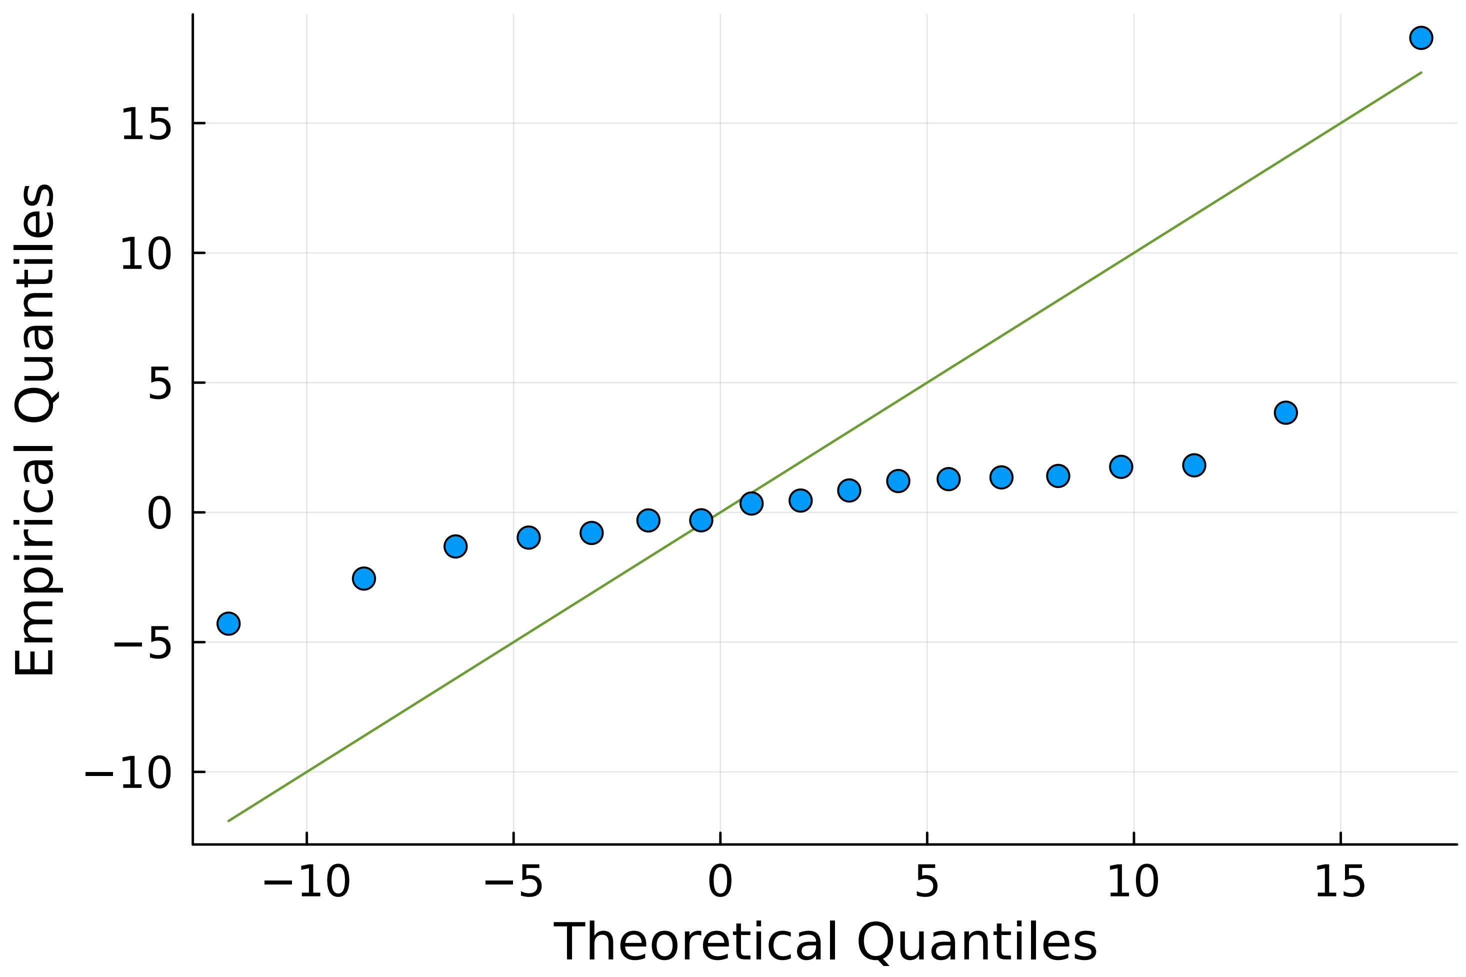

Autocorrelation

Another critical question is if the samples are **correlated** or **independent**. For a time series, this can be tested using autocorrelation (or cross-correlation for multiple variables).

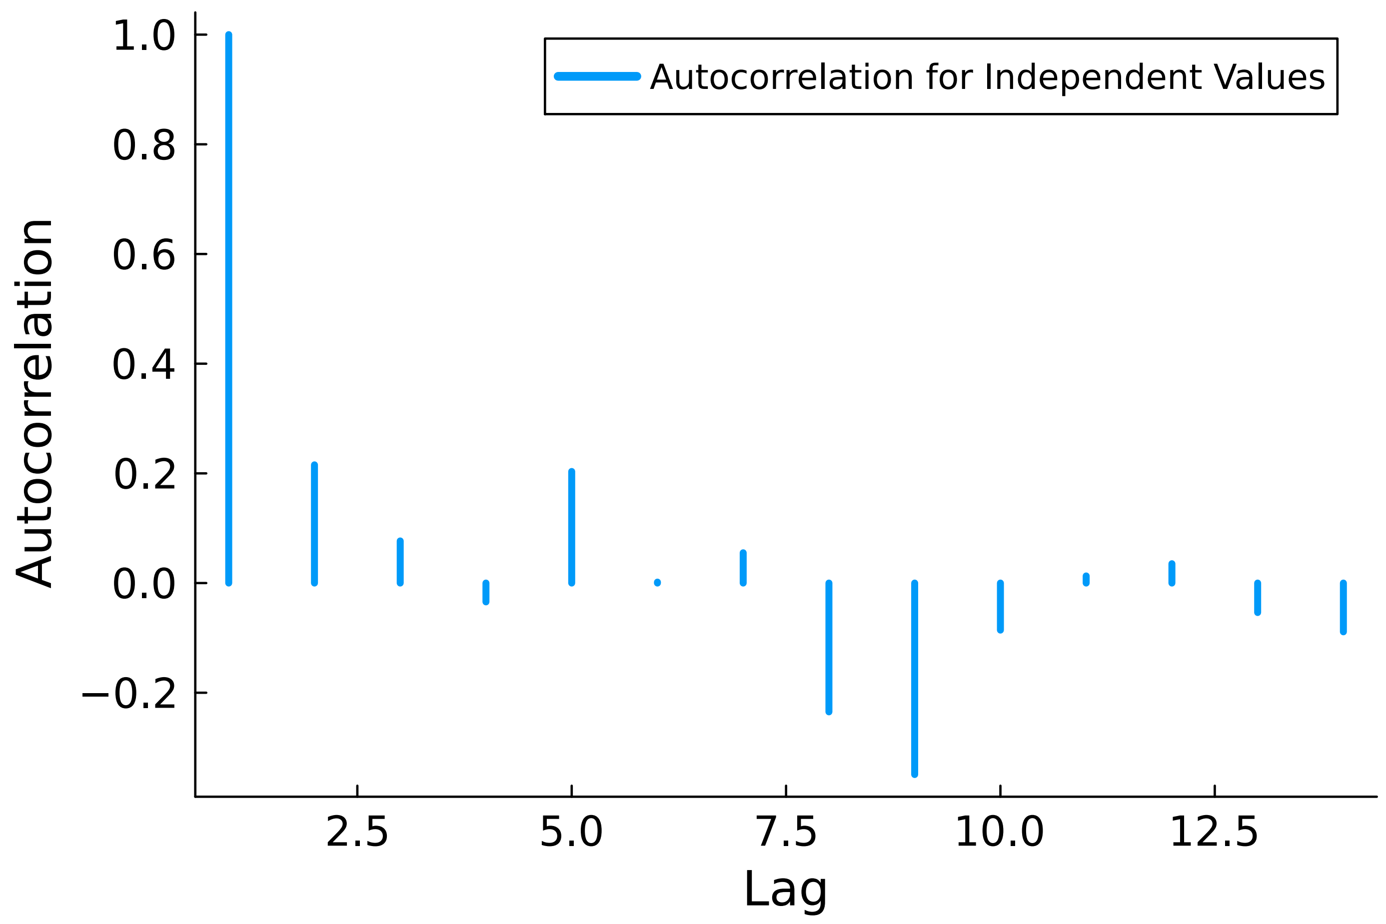

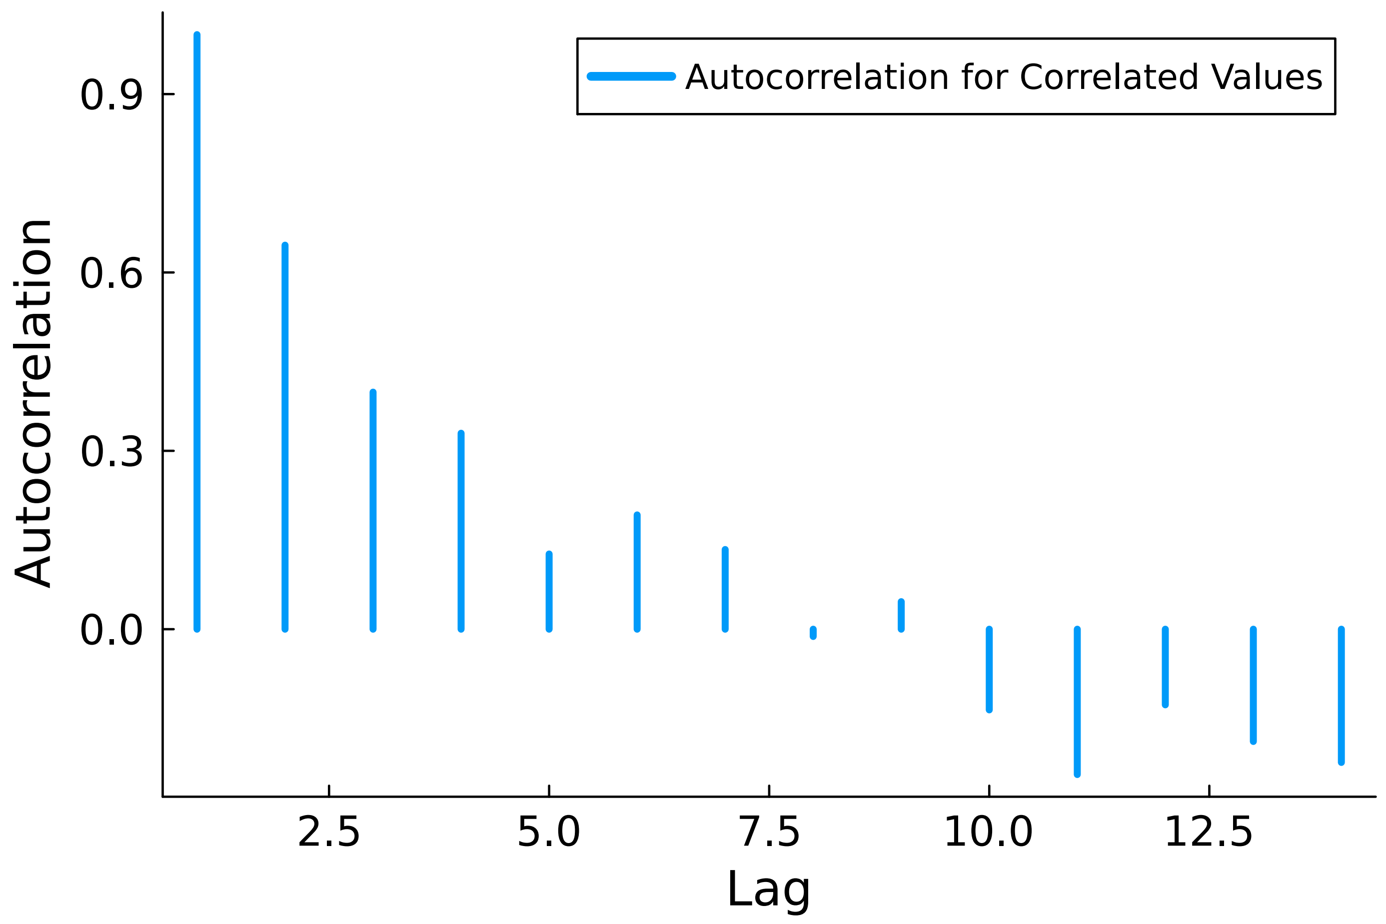

Key Takeaway

Specifying the probability model is important — getting this too wrong can bias resulting inferences and projections.

There's no black-box workflow for this: try exploring different methods, relying on domain knowledge, and looking at different specifications until you convince yourself something makes sense.

Monte Carlo

Monte Carlo Simulation

A common problem in risk/uncertainty analysis is *uncertainty propagation*: what is the impact of input uncertainties on system outcomes? The most basic way to approach this is through **Monte Carlo simulation**.

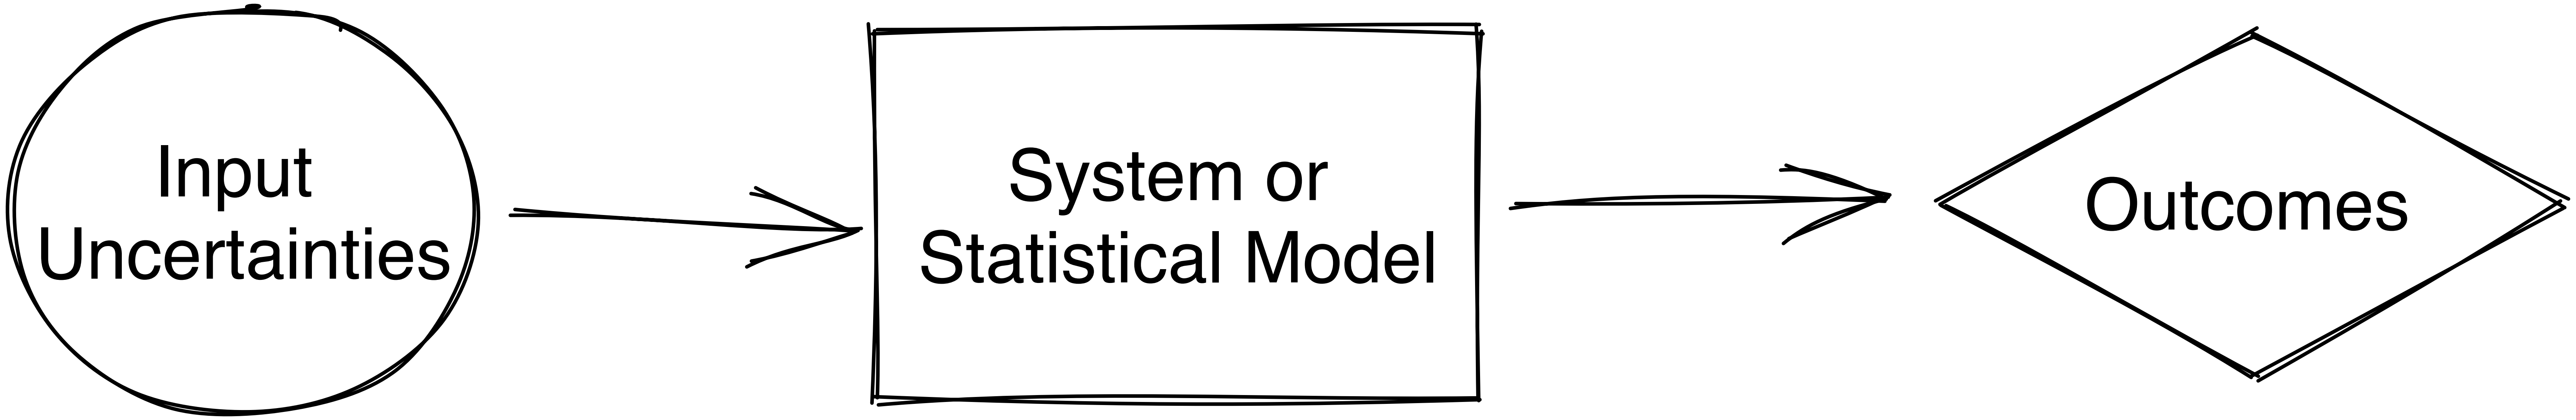

Monte Carlo Simulation

Monte Carlo simulation involves:

1. Sampling input(s) from probability distribution(s);
2. Simulating the quantity of interest;
3. *Aggregating the results* (if desired).

Monte Carlo Simulation

Monte Carlo simulation involves:

1. Sampling input(s) from probability distribution(s);
2. Simulating the quantity of interest;
3. *Aggregating the results* (if desired).

Note that steps 1 and 2 require the ability to **generate** data from the probability model (or we say that the model is **generative**). This is not always the case!

Monte Carlo Simulation

Monte Carlo is a very useful method for calculating complex and high-dimensional integrals (such as expected values), since an integral is an n-dimensional area:

1. Sample uniformly from the domain;
2. Compute how many samples are in the area of interest.

Monte Carlo (Formally)

We can formalize this common use of Monte Carlo as the computation of the expected value of a random quantity $f(Y)$, $Y \sim p$, over a domain D:

$$\mu = \mathbb{E}[f(Y)] = \int_D f(y)p(y)dy.$$

Monte Carlo (Formally)

Generate n independent and identically distributed values $Y_1, \ldots, Y_n$. Then the sample estimate is

$$\tilde{\mu} = \frac{1}{n} \sum_{i=1}^{n} f(Y_i)$$

The Law of Large Numbers

Monte Carlo works because of the **large of law numbers**:

If

1. Y is a random variable and its expectation exists and
2. $Y_1, \ldots, Y_n$ are independently and identically distributed

Then by the **strong law of large numbers**:

$$\tilde{\mu}_n \to \mu \text{ almost surely as } n \to \infty$$

Monte Carlo Estimators Are Unbiased

Notice that the sample mean $\tilde{\mu}_n$ is itself a random variable.

With some assumptions (the mean of Y exists and Y has finite variance), the expected Monte Carlo estimate is

$$\mathbb{E}[\tilde{\mu}_n] = \frac{1}{n}\sum_{i=1}^{n}\mathbb{E}[f(Y_i)] = \frac{1}{n}n\mu = \mu$$

This means that the Monte Carlo estimate is an *unbiased* estimate of the mean.

Ok, So That Seems Easy...

The *basic* Monte Carlo algorithm is straightforward: draw a large enough set of samples from your input distribution, simulate and/or compute your test statistic for each of those samples, and the sample value will necessarily converge to the population value.

However:

- Are your input distributions correctly specified (including correlations across inputs)?
- **How large is "large enough"?**

Monte Carlo Error

This raises a key question: how can we quantify the **standard error** of a Monte Carlo estimate?

The variance of this estimator is:

$$\tilde{\sigma}_n^2 = \mathrm{Var}\left(\tilde{\mu}_n\right) = \mathbb{E}\left(\left(\tilde{\mu}_n - \mu\right)^2\right) = \frac{\sigma_y^2}{n}$$

So the standard error $\tilde{\sigma_n}$ decreases approximately as $1/\sqrt{n}$ as n increases.

Monte Carlo Error

In other words, if we want to decrease the Monte Carlo error by 10x, we need 100x additional samples. **This is not an ideal method for high levels of accuracy.**

> *Monte Carlo is an extremely bad method. It should only be used when all alternative methods are worse.*
>
> *— Sokal, Monte Carlo Methods in Statistical Mechanics, 1996*

Monte Carlo Error

In other words, if we want to decrease the Monte Carlo error by 10x, we need 100x additional samples. **This is not an ideal method for high levels of accuracy.**

> *Monte Carlo is an extremely bad method. It should only be used when all alternative methods are worse.*
>
> *— Sokal, Monte Carlo Methods in Statistical Mechanics, 1996*

The thing is, though – for a lot of problems, all alternative methods *are* worse!

Reporting Monte Carlo Uncertainty

An **α-credible interval** for a Monte Carlo estimate is straightforward: compute an empirical interval containing α% of the Monte Carlo sample values (*e.g.* for a 95% credible interval, take the range between the 0.025 and 0.975 quantiles).

Monte Carlo Confidence Intervals

To estimate confidence intervals, we can rely on the variance estimate from before.

For "sufficiently large" sample sizes n, the **central limit theorem** says that the distribution of the error $|\tilde{\mu}_n - \mu|$ can be approximated by a normal distribution,

$$|\tilde{\mu}_n - \mu| \to \mathcal{N}\left(0, \frac{\sigma_y^2}{n}\right)$$

Monte Carlo Confidence Intervals

This means that we can construct confidence intervals using the inverse cumulative distribution function for the normal distribution.

The α-confidence interval is:

$$\tilde{\mu}_n \pm \Phi^{-1}\left(1 - \frac{\alpha}{2}\right)\frac{\sigma_Y}{\sqrt{n}}.$$

For example, the 95% CI is $\tilde{\mu}_n \pm 1.96\sqrt{\sigma_Y}/\sqrt{n}$.

Monte Carlo Confidence Intervals

Of course, we typically don't know σ_y. We can replace this with the sample standard deviation, though this will increase the uncertainty of the estimate.

But this gives us a sense of how many more samples we might need to get a more precise estimate.

A Dice Example (Cliche Alert!)

What is the probability of rolling 4 dice for a total of 19?

Let's solve this using Monte Carlo.

A Dice Example (Cliche Alert!)

What is the probability of rolling 4 dice for a total of 19?

Let's solve this using Monte Carlo.

- **Step 1**: Run n trials (say, 10,000) trials of 4 dice rolls each.
- **Step 2**: Compute the frequency of trials for which the sum is 19, *e.g.* compute the sample average of the indicator function

$$\frac{1}{n}\sum_{i=1}^{n}\mathbb{I}(\text{sum of 4 dice} = 19).$$

A Dice Example

How does this estimate evolve as we add more samples?

Note: the true value (given by the red line) is 4.32%.

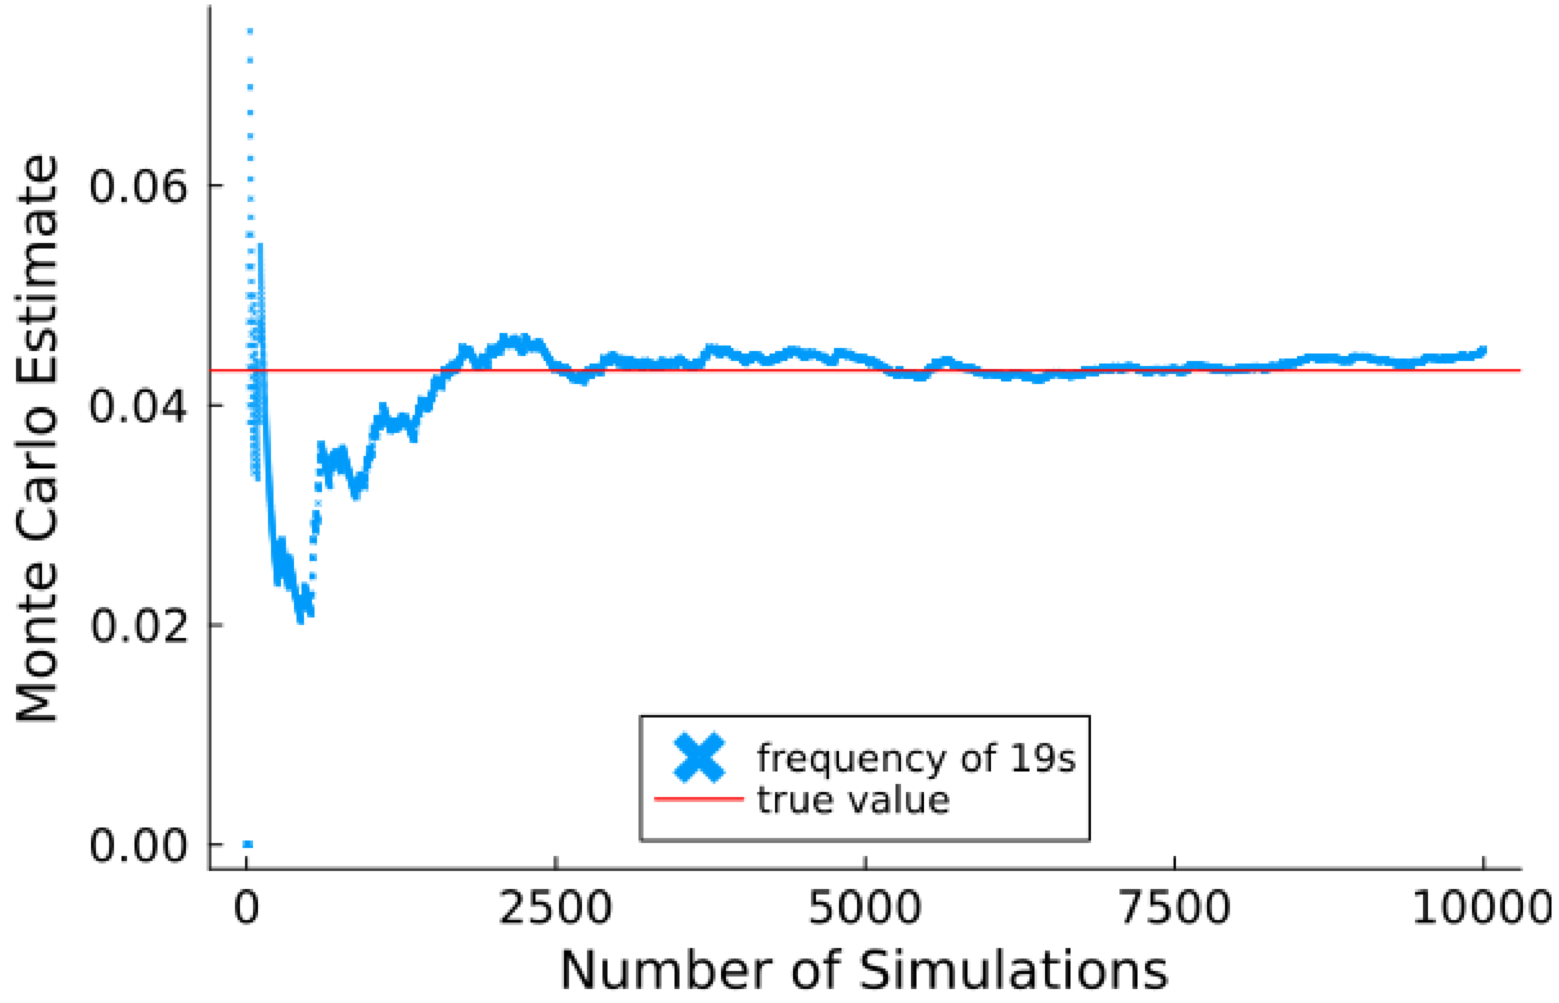

More Complex Monte Carlo

We won't spend too much more time here, but for more complex problems, the sample size needed to constrain the Monte Carlo error can be computationally burdensome.

This is typically addressed with more sophisticated sampling schemes which are designed to reduce the variance from random sampling, causing the estimate to converge faster.

- Importance sampling
- Quasi-random sampling (*e.g.* Sobol)

Key Takeaways (Monte Carlo)

- The basic Monte Carlo algorithm is a simple way to propagate uncertainties and compute approximate estimates of statistics, though its rate of convergence is poor.
- Can also be used for general simulation (which we will do later) and optimization.
- **Note**: Monte Carlo is a fundamentally *parametric* statistical approach, that is, it relies on the specification of the data-generation process, including all parameter values.
- What if we don't know these specifications *a priori*? This is the fundamental challenge of **uncertainty quantification**, which we will discuss more throughout this course.

Upcoming Schedule

Wednesday: Discuss Simpson (2021) and lab on testing for normality and Monte Carlo (featuring *The Price is Right*!).

Next Monday: Representing climate uncertainties and implications for risk management.