Probability Distributions & What they can Do for You!

Roadmap

So far, we've been slinging around normal distribution terminology casually. Let's formalize it, and make it useful for hypothesis testing!

1. Basic Probability Review
2. Other distributions: The world ain't Normal!
3. Our first mode of inference: P-Values

Probability!

Probability - The fraction of observations of an event given multiple repeated independent observations.

A Feeding Trial Example

Let's say you've offered 50 budworms a leaf to eat. 45 eat. $P(eats) = \frac{45}{50} = 0.9$

Now you offer 50 others a treated leaf. 10 eat. $P(eats) = \frac{10}{50} = 0.2$

Probability of NOT doing something

What is the probability of not eating if you are fed a treated leaf?

P(! eats) $= 1 - \frac{10}{50} = 0.8$

P(!A) = 1-P(A)

Probability of Exclusive Events

What if we offered our budworms both a treated and untreated leaf? 20 eat the control, 5 eat the treated leaf.

P(eats) $= \frac{20}{50} + \frac{5}{50} = 0.5$

P(A or B) = P(A) + P(B)

Two Events

We offer our budworms a leaf. 45 eat it. Then we offern them seconds. 20 of the original 45 each the second leaf.

P(eats twice) $= \frac{20}{50} = 0.4$

$= \frac{45}{50} * \frac{20}{45}$

P(A and B) = P(A)P(B)

Two Conditional Events

If we are interested in the probability of eating twice - i.e. the probability of eating a second time *given* that a budworm ate once, we phrase that somewhat differently.

$P(eats_2|eats_1)$

So, P(A given B) $= P(A|B)$
And thus, P(A and B) $= P(A)P(B|A)$

Distributions!

(when a point probabilty just ain't enough)

Frequency Distributions Make Intuitive Sense

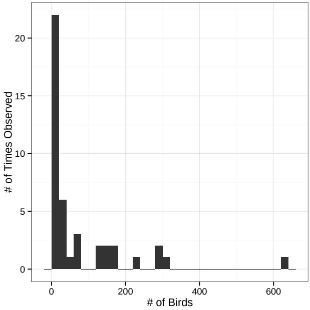

Frequencies Can be Turned Into Probabilities

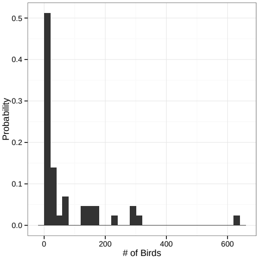

Just divide by total # of observations
But - we have binned observations...

Frequencies of Individual Observations

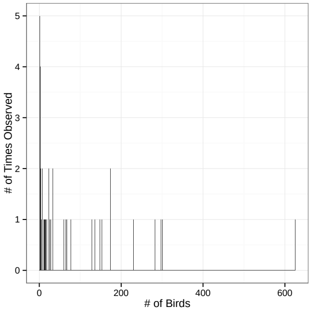

Can we turn these into probabilities?

Probabilities of Individual Measurements

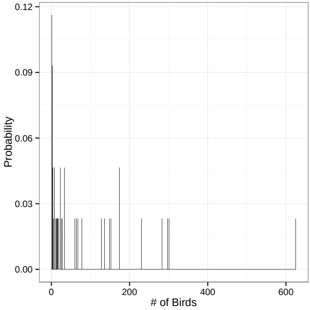

Many probabilities small, and what about the gaps?

Continuous Probability Distributions

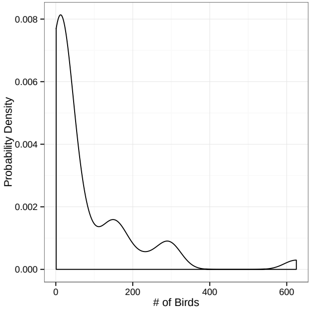

Any individual observation has a *probability density*.

Probability of a Range of Values

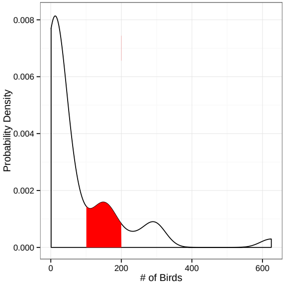

P(a=100) or P(a=101) or P(a=102)... = P(a=100) + P(a=101) + P(a=102)

$\int_{i=100}^{200} P(a=i)$

Probability as Integral Under the Curve

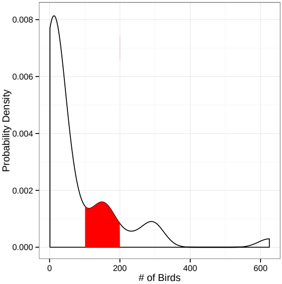

We obtain probabilities of observations between a range of values by integrating the distribution over selected values.

The Normal Distribution

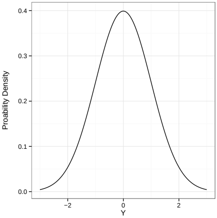

- Defined by it's mean and standard deviation.
- Y $\sim$N(μ, σ)
- Single mode
- Symmetric

67% of Values within 1 SD

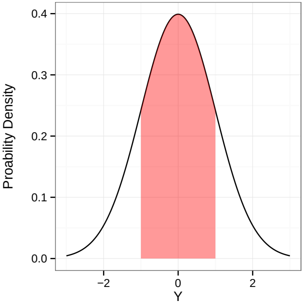

95% of Values within 2 (1.96) SD

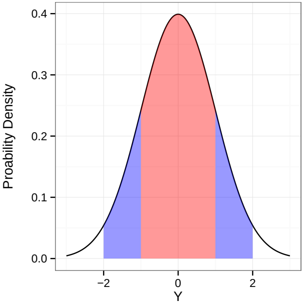

How to Get A Probability Density in R

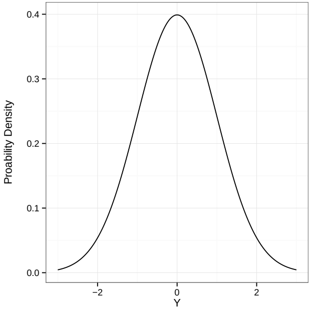


```
dnorm(Y, mean = 0, sd = 1)
```

The Probability of a Value or More Extreme Value

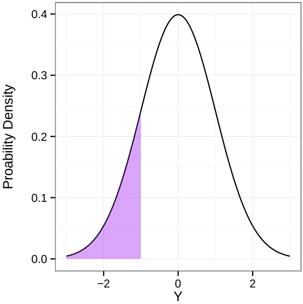


```
pnorm(Y, mean = 0, sd = 1)
```

The Cummulative Distribution/Quantile Function

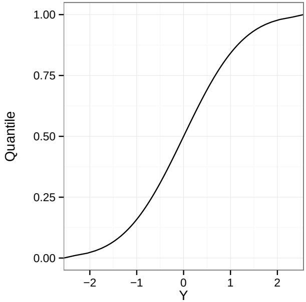


```
qnorm(p, mean = 0, sd = 1)
```

The Cummulative Distribution/Quantile Function

pnorm an qnorm are the inverse of one another

```
pnorm(-1)

# [1] 0.1587

qnorm(pnorm(-1))

# [1] -1

qnorm(0.025)

# [1] -1.96
```

The Lognormal Distribution

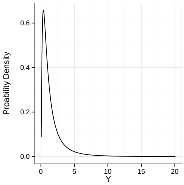

- An exponentiated normal
- Defined by the mean and standard deviation of its log.
- Y ~LN(μ_{log}, σ_{log})
- Generated by multiplicative processes

```
dlnorm(Y,
       meanlog=0,
       sdlog=1)
```

The Gamma Distribution

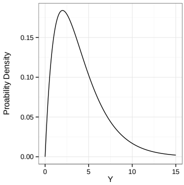

- Defined by number of events(shape) average time to an event (scale)
- Can also use rate (1/scale)
- Y ~G(shape, scale)
- Think of time spent waiting for a bus to arrive

```
dgamma(Y, shape = 2, scale = 2)
```

Waiting for more events

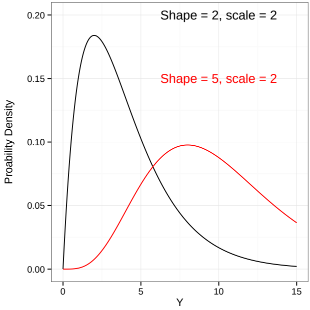

Longer average time per event

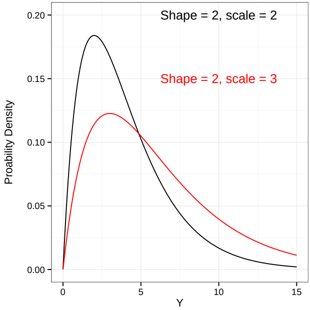

The Poisson Distribution

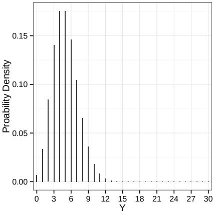

- Defined by λ - the mean and variance
- Y ~ P(lambda)

```
dpois(Y, lambda = 5)
```

When Lambda is Large, Approximately Normal

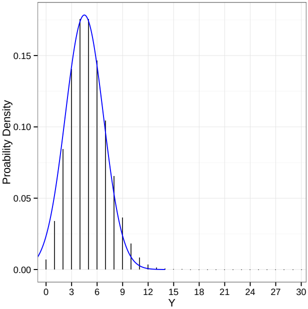

The Binomial Distribution

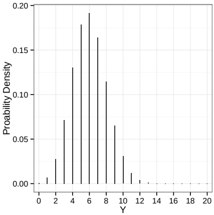

- Results from multiple coin flips
- Defined by size (# of flips) and prob (probability of heads)
- Y ~ B(size, prob)
- bounded by 0 and size

```
dpois(Y, size, prob)
```

Increasing Probability Shifts Distribution

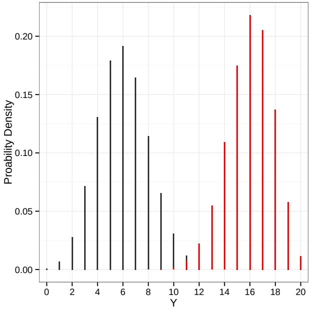

The Negative Binomial Distribution

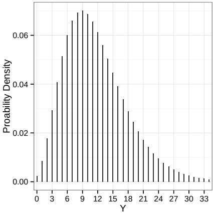

- Distribution of number of failures before n number of successes in k trials
- Or mean # of counts, μ, with an overdispersion parameter, size
- Y ~ NB(μ, size)

```
dnbinom(Y, mu, size)
```

Exercise

- Explore the distributions we have discussed
- Examine how changing parameters shifts the output of probability function
- Compare curves generated using density functions (e.g., dnorm) and large number of random draws (e.g. from rnorm)
- Overlay these in plots if you can (hist, lines, etc.)
- Challenge: graphically show integration under the different types of distribution curves (?polygon or ?geom_ribbon)

Hypothesis Testing

How Do we Derive Truth from Data?

Frequentist Inference: Correct conclusion drawn from repeated experiments

Bayesian Inference: Probability of belief that is constantly updated

Modes of Frequentist Inference

Null Hypothesis Tests: Falsify a null hypothesis
Likelihood/Information Theoretic: Evaluate weight of evidence

Inductive v. Deductive Reasoning

Deductive Inference: A larger theory is used to devise many small tests. NHT.

Inductive Inference: Small pieces of evidence are used to shape a larger theory. Likelihood.

Null Hypothesis Tests & Popper

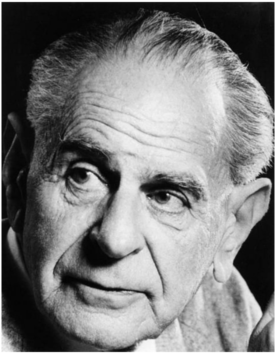

Falsification of hypotheses is key!

A theory should be considered scientific if, and only if, it is falsifiable.

Deductive Reasoning and Null Hypothesis Tests

A null hypothesis is a default condition that we can attempt to falsify.

Common Uses of Null Hypothesis Tests

- Ho: Two groups are the same
- Ho: An estimated parameter is not different from 0
- Ho: The slopes of two lines are the same
- Etc...

So, what conclusions can we draw if we reject the null?

Ho and Ha

There are often many alternate hypotheses. Rejection of the null does not imply acceptance of any single alternative hypothesis.

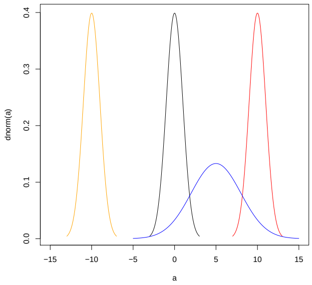

Null Distributions

Null hypotheses are associated with null statistical distributions.

For example, if Ho states that a value is normally distributed, but is not different from 0, the null distribtion is centered on 0 with some standard deviation.

Null Distributions

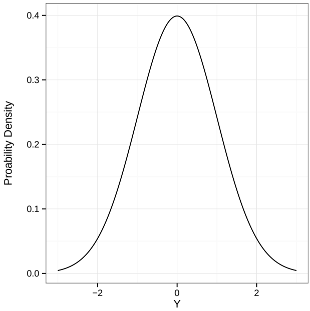

The P Value

P-value: The Probability of making an observation or more extreme observation given that the null hypothesis is true.

R. A. Fisher

Evaluation of a Test Statistic

We can use our data to calculate a test statistic that maps to a value of the null distribution. We can then calculate the probability of observing our data, or of observing data even more extreme, given that the null hypothesis is true.

$$P(X \leq Data|H_0)$$

Evaluation of a Test Statistic

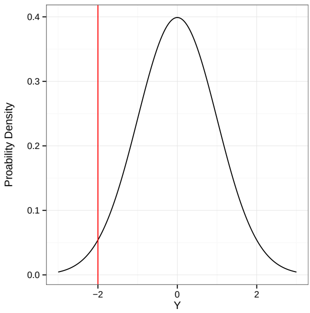

The P Value

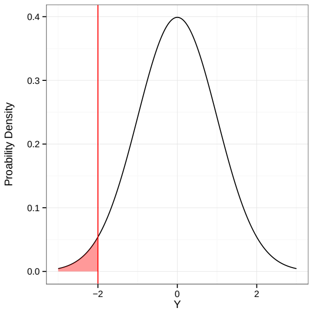

p=0.0227, Note - this is a one-tailed test!

1-Tailed v. 2-Tailed Tests

1-Tailed Test: We are explicit about whether Ha implies that our sample is greater than or less than our null value.

$$P(X \leq Data|H_0) \text{ (1-tailed)}$$

2-Tailed Test: We are make no assumption about the sign or direction of our alternative hypotheses.

$$P(X \leq Data|H_0) + P(X \geq Data|H_0) \text{ (2-tailed)}$$

Two-Tailed P Value

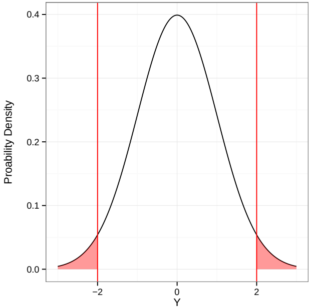

p=0.0454 from `pnorm(-2)*2`

When should you use a 1-Tailed Test?

What does 0.454 mean?

There is a 4.54% chance of obtaining the observed data, or more extreme data, given that the null hypothesis is true.

If you chose to reject the null, you have a 1 in 22 chance of being wrong.
How comfortable are you with rejecting the null?

Note: rejecting the null $\neq$ accepting a specific alternative

Exercise: Evaluate Support for Null Hypothesis

- Typically, the number of warts on a toad is Poisson distributed with a λ of 54
- You survey a lake suspected to contain high PAH levels. You pick up a toad, and it has 40 warts.
- What is your null hypothesis?
- What is the probability of making this observation, given your null?
- Challenge: How does your p value change with # of warts, say, from 1 to 108 warts?

Exercise: Evaluate Support for Null Hypothesis

```r
2*ppois(40, 54)

# [1] 0.05755

#OR!
p<-0
for(i in 1:40){
  p<-p+dpois(i, 54)
}
p*2

# [1] 0.05755
```

Exercise: Evaluate Support for Null Hypothesis

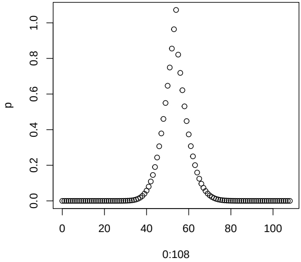

Exercise: Evaluate Support for Null Hypothesis

```
p<-0
for(i in 0:54){
  p[i+1]<-2*ppois(i, 54)
}

for(i in 55:108){
  p[i+1]<-2*ppois(i, 54, lower.tail=F)
}

plot(0:108, p)
```