

Continuous Distributions

Homework Problems

Homework Solutions

1. Continuous Distributions
 - 1.1. Problem: Some Continuous RV
2. The Uniform
3. The Normal Distribution
 - 3.1. Problem: Four Normals
 - 3.2. Problem: Standard Normal
4. The Cumulative Distribution Function
 - 4.1. Problem: Normal Return
 - 4.2. Problem: Uniform CDF
 - 4.3. Problem: Inverse CDF
5. The Long Run: Probabilities and Means
 - 5.1. Problem: Mean of IID Normal

1. Continuous Distributions

Suppose we have a machine that cuts cloth.
When pieces are cut, there are remnants.

We believe that the length of a remnant could be anything between 0 and .5 inches and, any value in the interval is equally likely.

The machine is about to cut, leaving a new remnant.

The length of the remnant is a number we are unsure about, so it is a random variable.

We can't list out all the possible values between 0 and .5!!!

In this case, we need a new way to talk about probability.

Instead of specifying the probability of particular values, we will give the probability of intervals.

Instead of

$$P(X = x)$$


we will talk about,


$$P(a < X < b).$$

One convenient way to specify the probability of any interval is with the probability density function (pdf).

The probability of an interval is the area under the pdf.

In this example, intervals near 0 are more likely.






For this random variable the probability that it is in the interval $[0,2]$ is .477. (47.7 percent of the time it will fall in this interval).

Note The area under the entire curve must be 1 (Why?)


Here is another p.d.f:



Most of the probability is concentrated in 1 to 15, but you could get a value much bigger. This kind of distribution is called skewed to the right.

1.1. Problem: Some Continuous RV

Here is the pdf of some continuous random variable.
(Note: the pdf = 0 for all $x < 0$).



- (a) What is the probability of being less than 0?
- (b) What is the probability of being greater than 30?
- (c) What is the probability of being less than 5.3?
(i) .1 (ii) .95 (iii) .5 (iv) 0
- (d) The probability of being between 5.3 and 10.6 is?
(i) .1 (ii) .95 (iii) .4 (iv) 0

Solution

(a) 0

(b) 0

(c) (iii) = .5

(d) (iii) = .4

2. The Uniform

Let's go back to our "remnant" example.


Any value between 0 and .5 is equally likely.

If X is the length of a remnant what is its pdf?

The pdf is flat in the interval $(0, .5)$.

Its height must be 2 so that the total area is 1.

Outside of $(0, .5)$, the pdf is 0.




Technically, any particular value has probability 0!
But we still express the idea that they are equally likely in $(0, .5)$!

In general, we have a *family* of uniform distributions describing the situation where any value in the interval (a,b) is equally likely.

$$f(x) = \begin{cases} \frac{1}{b-a} & a < x < b \\ 0 & \text{otherwise} \end{cases}$$

We write $X \sim U(a,b)$.




3. The Normal Distribution

This pdf describes the *standard normal distribution*.

We often use Z to denote the RV which has this pdf.

Note:
any value in
 $(-\infty, \infty)$
is "possible".



The standard normal does not seem that useful.


How often do we want say the number will likely be between -2 and 2?

We get a *family* of normal distributions by moving the bell curve around with the parameter μ and stretching/shrinking it with the parameter σ .

$X \sim N(\mu, \sigma^2)$ means X has this pdf:

$$P(\mu - 2\sigma < X < \mu + 2\sigma) = .95$$


$$P(\mu - \sigma < X < \mu + \sigma) = .68$$



The normal family has two parameters

μ : where the curve is centered

σ : how spread out the curve is



Z, X, and Y are all "normally distributed".

Example:


<http://www.britannica.com/EBchecked/topic/289766/human-intelligence/13355/The-distribution-of-IQ-scores>

Intelligence test scores follow an approximately normal distribution, meaning that most people score near the middle of the distribution of scores and that scores drop off fairly rapidly in frequency as one moves in either direction from the center. For example, on the IQ scale, about 2 out of 3 scores fall between 85 and 115, and about 19 out of 20 scores fall between 70 and 130. Put another way, only 1 out of 20 scores differs from the average IQ (100) by more than 30 points.

$$1/20 = .05.$$

So, $\mu = 100$ and $\sigma = 15$.

$IQ \sim N(100, 225)$. ($225 = 15^2$)



Note:

$Z \sim N(0, 1)$ is the *standard normal*.

Later on, we will see that it is useful.

Note:

Later on we will see that for $X \sim N(\mu, \sigma^2)$,

$E(X) = \mu$ and $Var(X) = \sigma^2$.

That is, μ is the mean or expected value of X , σ^2 is the variance of X , and σ is the standard deviation.

Note:

It is really, $P(\mu - 1.96 \sigma < X < \mu + 1.96 \sigma) = .95$.

but we will just say probability of $\mu \pm 2 \sigma = .95$.

Note:

μ : *what is going to happen ??*





σ : *how sure are you ??*

3.1. Problem: Four Normals

At right are pdfs for X_i , $i = 1, 2, 3, 4$.

Each X_i is normal with μ equal to one of $-1, 1, 2,$ and 0 .

σ is either 1 or $.5$.



(a) What μ and σ for X_1 ?

(b) What μ and σ for X_2 ?

(c) What μ and σ for X_3 ?

(d) What μ and σ for X_4 ?

Solution

(a) $(\mu, \sigma^2) = (-1, .25)$.

(b) $(\mu, \sigma^2) = (1, 1)$.

(c) $(\mu, \sigma^2) = (2, 1)$.

(d) $(\mu, \sigma^2) = (0, .25)$.

3.2. Problem: Standard Normal

Suppose Z has the standard normal distribution ($Z \sim N(0, 1)$).

- (a) What is $P(Z > 1.96)$?
- (b) What is $P(Z > -1.96)$?
- (c) What is $P(0 < Z < 1.96)$?
- (d) What is $P(-1 < Z < 1.96)$?
- (e) What is $P(-1 < Z < 1)$?
- (f) What is $P(-2 < Z < 2)$?

Solution

(a) $.05/2 = .025$.

(b) $1-.025 = .975$.

(c) $.5-.025=.475$.

(d) $(.68)/2 + .475 = .815$. (if you use the cdf function, you get a more accurate answer of .8163.)

(e) .68

(f) .9545 with cdf, or about .95.

4. The Cumulative Distribution Function

The c.d.f. (**cumulative distribution function**) is just another way (besides the p.d.f.) to specify the probability of intervals.

For a random variable X the c.d.f., which we denote by F (we used f for the p.d.f.), is defined by

$$F_X(x) = P(X \leq x)$$


The c.d.f. is handy for computing the probabilities of intervals.


$$\begin{aligned}P(a < X \leq b) &= P(X \leq b) - P(X \leq a) \\ &= F_X(b) - F_X(a)\end{aligned}$$

Example

For Z (standard normal),
we have:

$$\begin{aligned}P((-1,1)) &= F(1) - F(-1) \\ &= .84 - .16 = .68\end{aligned}$$






The probability of an interval is the *jump* in the c.d.f. over that interval.

Note: for x big enough, $F(x)$ must get close to 1.
for x small enough, $F(x)$ must get close to 0.

Example

Let R denote the return on our portfolio next month. We do not know what R will be. Let us assume we can describe what we think it will be by:

$$R \sim N(.01, .04^2)$$



What is the probability of a negative return?

In excel we use:

=NORMDIST(0,0.01,0.04,TRUE)

$F_X(0), X \sim N(.01, .04^2)$

And then the cell will be: .4013

$$P(R < 0) = \text{cdf}(0) = .4013$$

What is the probability of a return between 0 and .05?

=NORMDIST(.05,0.01,0.04,TRUE) = .8413

$$P(0 < R < .05) = .84 - .4 = .44$$

$F_X(.05), X \sim N(.01, .04^2)$

4.1. Problem: Normal Return

Suppose R is a RV representing returns on an asset next month and $R \sim N(.05, .01)$.

(a)

Give an interval such that R has a 95% chance of being in that interval?

(b)

What is $\text{Prob}(R > .05)$?

(c)

What is $\text{Prob}(R > .25)$?

(d)

What is $\text{Prob}(-.05 < R < .15)$?

(e)

What is $\text{Prob}(R < .05)$?

(f)

What is $\text{Prob}(R < .025)$? (for this one you will need the cdf!)

In excel:

```
=normdist(.025,.05,.1,TRUE)
```

In R:

```
> pnorm(.025,mean=.05,sd=.1)
```

```
[1] 0.4012937
```

(g)

I gave you the answer to (f)!!

What is $P(0 < R < .025)$?

(h)

Plot normal pdf (probability density function) for $R \sim N(.05, .01)$.

In Excel:

- ▶ create a column of values from $-.35$ to $.45$ (($\mu - 4\sigma$) to ($\mu + 4\sigma$)).
You can use /Home/Editing/Fill/Series.
(the downarrow thing, put $-.35$ in a cell and click on it, go to fill, check columns, then use $.02$ for step and $.45$ for “Stop value”).
- ▶ copy the formula `=normdist(a2,.05,.1,FALSE)` into the cells beside the column of values you just created (I assume the values started in cell a2).
- ▶ Now you will have two columns side by side. click on a cell in the two columns, click Ctrl-a to select all the values, /Insert/Scatter. Then “click around to format the plot the way you like.

In R:

```
mu=.05
sigma=.1
rvals = seq(from=mu-4*sigma,to=mu+4*sigma,length.out=100) #grid of points
pdfv = dnorm(rvals,mean=mu,sd=sigma) #normal pdf evaluated on grid
plot(rvals,pdfv,type="l",col="blue",xlab="return",ylab="")
```

Solution

Note: $\sigma = .1$, $\mu = .05$.

(a) $\mu \pm 2\sigma$!! .95.

(b) $P(R > \mu) = .5$.

(c) $P(R > \mu + 2\sigma) = .025$.

(d) $\mu \pm \sigma$!! .68.

(e) $P(R < \mu) = .5$.

(f)

```
> pnorm(.025, .05, .1)
```

```
[1] 0.4012937
```

4.2. Problem: Uniform CDF


Suppose $X \sim U(0,1)$ (uniform on the interval $(0,1)$).

Plot the cdf (cumulative distribution function) of X from -0.5 to 1.5 .

Solution

For x in $(0,1)$,

$$F(x) = P(X \leq x) = \text{width} \times \text{height} = x * 1 = x.$$



4.3. Problem: Inverse CDF

In excel evaluate

=norminv(.5,0,1)

=norminv(.025,0,1)

=norminv(.16,0,1).

Can you figure out what norminv does?

Note: in R, instead of norminv use:

```
> qnorm(.5,0,1)
```

```
> qnorm(.025,0,1)
```

```
> qnorm(.16,0,1)
```

For $R \sim N(.05, .01)$, what is the value of r such that $P(R < r) = .1$?

Solution

Given a value, the CDF tells us the probability that our RV will end up being less than that value.

The *inverse CDF* goes the other way.

Given a probability, it tells us the value such that the chance of being less than that value equals the given probability.

```
> qnorm(.1, .05, .1)
[1] -0.07815516
> pnorm(-0.07815516, .05, .1)
[1] 0.09999999
```

5. The Long Run: Probabilities and Means

I tossed two coins 2000 times.

Each time, I recorded the number of heads.

```
2 1 1 0 2 2 0 0 2 2 1 2 0 1 0 0 1 1 1 1 0 1 2 2 1
2 1 2 1 1 2 1 1 2 2 0 0 1 1 1 1 2 1 0 0 1 0 2 1 2
2 1 1 2 0 2 0 1 0 0 0 2 1 2 1 0 2 2 2 1 1 1 2 1 1
1 2 0 2 2 1 2 0 2 1 1 1 0 2 1 1 0 1 1 2 0 1 1 0 0
1 1 2 2 1 1 1 0 1 2 2 0 1 1 0 1 2 2 2 2 1 1 2 0 0
0 0 2 2 1 2 1 1 1 2 1 1 2 0 0 1 0 0 1 1 2 0 2 1 1
1 0 1 1 2 0 0 2 1 1 2 1 1 2 1 1 1 1 1 1 0 2 1 2 0
2 0 0 2 0 1 1 1 0 1 1 1 2 2 1 1 1 2 0 1 0 1 0 0 1
0 1 0 1 1 1 1 1 1 1 1 2 1 1 2 0 2 0 2 1 0 2 2 0 0
1 1 2 1 2 1 0 0 1 1 0 1 0 1 0 1 0 1 2 1 1 1 0 1 1
```

.....

.....

```
0 0 1 0 2 1 0 2 0 0 1 2 1 1 1 1 2 2 2 2 2 1 1 2 1
2 0 2 1 1 1 0 0 2 1 1 1 1 1 1 2 2 2 1 0 1 0 1 1 1
1 0 1 1 2 1 2 2 2 2 1 2 1 0 1 2 2 1 2 2 1 1 0 1 2
1 1 1 1 1 1 1 2 0 1 0 2 0 1 2 1 1 2 1 1 2 1 2 2 2
0 2 1 2 0 2 1 2 0 1 1 2 2 2 2 0 1 2 1 1 1 1 1 0 2
0 0 1 0 0 0 1 1 2 2 1 2 1 1 0 1 0 1 0 0 1 1 1 0 2
1 2 0 0 1 2 0 1 0 1 2 1 1 1 0 1 1 2 1 1 2 1 2 1 2
1 1 2 0 0 0 1 0 2 1 1 0 2 1 2 2 2 2 2 0 1 2 1 1 1
0 0 1 1 1 2 1 0 1 1 0 1 1 0 1 0 2 0 0 2 1 2 0 1 1
2 2 0 0 0 1 1 1 1 2 2 1 0 0 1 1 0 2 0 1 1 1 0 0 1
2 2 1 1 1 1 1 0 1 1 0 1 0 1 1 1 2 1 0 1 2 2 1 1 2
```



We can think of the numbers as IID draws Y_i .
Each Y_i has the same distribution as

Y :

y	$Pr(Y = y)$
0	.25
1	.5
2	.25

What fraction of the numbers are equal to 1?

This is the running proportion of times we get 1.



For each toss #, we have the fraction of tosses up to and including the toss that resulted in 1 head.


Let n_1 denote the number of times we got a 1 out of n tosses..

Then, *in the long run* (that is, for large n)

$$\frac{n_1}{n} \approx P(Y = 1) = .5$$

We can think of the probability of an outcome as the long run frequency of that outcome in a sequence of IID draws.

10 times, I tossed 2 coins 2,000 times.
Proportion of times I got 1 head.



Note:

We have already used this reasoning !!!

We modeled defects as $Y_i \sim \text{Bernoulli}(p)$.

We can interpret p as the long run proportion of defects if we run the process forever.

Since, **given the model**, the sample proportion would be like the “true” p , in the long run, we use it to *estimate* p .

Note: We can always use the sample quantity to summarize the data.

We had 18% defectives.

In the context of our model we can also use the sample quantity to estimate the parameter of our model.

*Our estimate of the parameter p
(the true, long run proportion of defectives)
is the sample proportion .18.*

Some folks also call p the “population” quantity.

Notation:

When we use the sample proportion as an estimate of a true (population p) we often use the “hat” notation:



Our estimate of the true p turned out to be $\hat{p} = .18$.

Long Run Proportion and Probabilities, Continuous Distributions

This works for continuous distributions as well, *except* (of course) we can only talk about the probability of intervals.



The long run proportion of IID draws that land in an interval is the probability of that interval.

IID draws $Y_i \sim U(0,1)$.



Proportion of draws in $(.2, .8)$.


IID draws $Y_i \sim N(10, 3^2)$.



Proportion of draws in $(7, 13)$.

Long Run Average of IID Draws

Here are my 2,000 tosses of two coins again.




What is the average of the numbers of these numbers?

Let n_0 be the number of 0's, n_1 be the number of 1's, and n_2 be the number of 2's.

$$\begin{aligned}\text{average} &= \frac{\text{sum}}{2000} \\ &= \frac{n_0 \times 0 + n_1 \times 1 + n_2 \times 2}{n} \\ &= \frac{n_0}{n} \times 0 + \frac{n_1}{n} \times 1 + \frac{n_2}{n} \times 2 \\ &\approx P(Y = 0) \times 0 + P(Y = 1) \times 1 + P(Y = 2) \times 2 \\ &= E(Y) \\ &= .25 \times 0 + .5 \times 1 + .25 \times 2 = 1.\end{aligned}$$

The long run average of IID draws is equal the the expected value under the distribution you are drawing from.



$Y_i \sim Y$ IID, $i = 1, 2, \dots, n \Rightarrow \bar{Y} \approx E(Y)$ for n "big".

Expected Values for Continuous Random Variables

If Y is a continuous random variable with pdf $f(y)$, then,

$$E(Y) = \int y f(y) dy$$

If you remember your calculus, this is analogous to what we did for discrete random variable.

If you don't remember your calculus *that is fine*.

We have a general intuition for what the expected value mean that works the same for discrete and continuous random variables.

For are random variable Y (discrete or continuous),

$E(Y)$ is

- ▶ the probability weighted average of the possible outcomes.
- ▶ the long run average of IID draws from the distribution of Y .

Suppose $Y \sim U(0, 1)$, what is $E(Y)$?

Suppose $Y \sim N(10, 3^2)$, what is $E(Y)$?

Suppose $Y \sim N(\mu, \sigma^2)$, what is $E(Y)$?

Variance for Continuous Random Variables

If Y is a continuous random variable with pdf $f(y)$, then,

$$\text{Var}(Y) = \int (y - \mu)^2 f(y) dy$$

We also write σ_Y^2 for the variance of Y and we have:

$$\text{sd}(Y) = \sigma_Y = \sqrt{\sigma_Y^2}.$$

Note:

Not obvious but true:

$$Y \sim N(\mu, \sigma^2) \rightarrow \text{Var}(Y) = \sigma^2$$

Note:

- ▶ the variance (or sd) measures how *spread out* the distribution is.
- ▶ the sample variance of iid draws gets close to the true variance as the number of draws get large.

5.1. Problem: Mean of IID Normal

Get 1,000 IID draws from the $N(100, 25)$ distribution.

Get the sample mean.

How does it compare to the “true” mean of 100?

In R:

```
set.seed(99) #99 for Gretzky
x = rnorm(1000,100,5)
print(mean(x))
##plot the data
plot(x)
abline(h=100,col="blue") #true mean
abline(h=mean(x),col="red") #sample mean
```