Stochastic Components of Models

Gov 2001 Section

February 5, 2014

Outline

1. Replication Paper and other logistics

2. Data Generation Processes and Probability Distributions

3. Discrete Distributions

4. Continuous Distributions

5. Simulating from Distributions

6. Distribution Transformations

Replication Paper

- Read *Publication, Publication* on Gary's website.
- Find a partner to coauthor with.
- Find a set of papers you would be interested in replicating.
  1. Recently published (in the last two years).
  2. From a top journal in your field.

     Think *American Political Science Review* or *American Economic Review*, not the Nordic Council for Reindeer Husbandry Research's *Rangifer*. Unless your field is reindeer husbandry.
  3. Use methods at least as sophisticated as in this class (more than just basic OLS).
- Have a classmate (somebody in your study group, preferably) approve your choice of article. They should make sure it fits the criteria listed in *Publication, Publication*.
- Each set of partners should email us (Gary, Stephen, and Solé) with a PDF of your article, a brief (2-3 sentence) explanation for why you picked it. And the name of the person who checked that it met the requirements in *Publication, Publication*.

Canvas

- Everybody who is
  1. registered for the class through Harvard
  2. cross-registered from MIT or a different Harvard department
  3. registered through the Extension School
  4. formally auditing the course through the Extension School

  should have access to Canvas by now.

- Informal auditors (people who didn't register for the class through the Extension School) can only get access to Canvas if they have a current Harvard affiliation (student, faculty, visiting scholar, etc.) and Harvard email address

Canvas

- Use the discussion board on Canvas. You'll get an answer quicker than if you email us. Please don't paste big chunks of code though. You don't want people copying your work.

- Everyone should change their notification preferences in Canvas. This will make sure you get an email when somebody responds to a post that you started or commented on.
  First you have to make sure your email address is linked to Canvas. Go to 'Settings' in the top right, then on the far right under 'Ways to Contact' you should see your email address. If it's not there, be sure to add it.
  Then change your notification settings by by clicking 'Settings' in the top right, and then 'Notifications' on the left, and choosing 'Notify me right away' for 'Discussion' and 'Discussion Post'

This week's homework

- By later tonight there will be two assignments on the Quizzes section on Canvas
    1. The second problem set, which you'll complete exactly as you did last week
    2. An assessment problem, which you must complete **independently** and can only submit one time to Canvas
- We won't answer any questions about R code or anything on the assessment problem. If you have a clarifying question, email all three of us and we'll post an answer on Canvas if we think we need to clarify something about the question.
- Comment about multiple submissions in Canvas.

Outline

Stochastic Components of Models

How do we build or select a statistical model?

Imagine a friend comes to you with a bunch of data, and they want you to help build a model that helps predict some outcome of interest.
What are the first questions you should ask them?

1. What is the dependent variable?

2. What was the data generation process that created that dependent variable?

Fundamental Uncertainty and Stochastic Models

"If we played them ten times, they might win nine. But not this game. Not tonight."

- Coach Herb Brooks in *Miracle*

Or put another way:

"If **Y** is whether the Soviet team beats us, then $Y \sim Bern(0.9)$. But

$y_{thisgame} \neq 1$. $y_{tonight} = 0$."

- Coach Herb Brooks in *Miracle*

Even the most certain things have fundamental uncertainty around them.

We use the stochastic component of our models to capture this fact.

So why become familiar with probability distributions?

Several reasons:

- You can pick the probability distribution that corresponds with the data generation process of your dependent variable.
- You can fit models to a variety of data. Not just Normal data!
- You can help your friends build a good statistical model for predictive or descriptive inference

What do we have to know about probability distributions in order to apply them to a particular problem?

- You have to know the "stories" behind them
- You know to pick a distribution which corresponds with the DGP for your dependent variable
- You have to understand the assumptions you're making them you pick that distribution

Outline

1. Replication Paper and other logistics

2. Data Generation Processes and Probability Distributions

3. Discrete Distributions

4. Continuous Distributions

5. Simulating from Distributions

6. Distribution Transformations

The Bernoulli Distribution

- **The story:** flipping a *single* coin
- The Bernoulli distribution has one parameter, π, which is the probability of "success".
- If $Y \sim Bern(\pi)$, then $y = 1$ with success probability π and $y = 0$ with failure probability $1 - \pi$.
- Ideal for modeling **one-time** yes/no (or success/failure) events.
- Other examples:
    - one voter voting yes/no
    - one person being either a man/woman

The Bernoulli Distribution

$Y \sim \text{Bernoulli}(\pi)$

y can take on a value of either 0 or 1. Nothing else.

probability of success: $\pi \in [0, 1]$

$p(y|\pi) = \pi^y (1 - \pi)^{(1-y)}$

$E(Y) = \pi$

$\text{Var}(Y) = \pi(1 - \pi)$

```
rbinom(100, size = 1, prob = .7)
```

The Binomial Distribution

- **The story:** flipping a coin a bunch of times and counting how many times it came up heads
- The Binomial distribution is the total of a bunch of Bernoulli trials.
- Two parameters: probability of "success", π, and number of trials, n Generally you *must* know the number of trials that generated your data. If you don't, or if there's a huge number of trials, you might want to use a different distribution.
- Examples:
  - You flip a coin three times and count the total number of heads you got. (The order doesn't matter.)
  - The number of women in a group of 10 Harvard students
  - The number of rainy days in the seven-day week

The Binomial Distribution

Histogram of Binomial(20,.3)

$Y \sim \text{Binomial}(n, \pi)$

$y = 0, 1, \ldots, n$

number of trials: $n \in \{1, 2, \ldots\}$

probability of success: $\pi \in [0, 1]$

$p(y|\pi) = \binom{n}{y} \pi^y (1 - \pi)^{(n-y)}$
$E(Y) = n\pi$

$\text{Var}(Y) = n\pi(1 - \pi)$

```
rbinom(100, size = 5,
prob = .7)
```

The Multinomial Distribution

- **The story:** rolling a dice (even or unevenly weighted) a bunch of times and counting how many times each number comes up
- Generalization of the binomial, which is just a two-dimensional multinomial
- Two parameters: number of trials, n, and a vector of probabilities of each of the k outcomes, $\mathbf{p} = \{p_1, p_2, ...p_k\}$
- Multinomial assumes that you have mutually exclusive outcomes.

- Examples:
    - you toss a die 15 times and get outcomes 1-6
    - election vote totals where there's 3+ candidates to choose from

The Multinomial Distribution

$Y \sim$ Multinomial$(n, \pi_1, \ldots, \pi_k)$

$y_j = 0, 1, \ldots, n$; where $\sum_{j=1}^{k} y_j = n$ must be true

number of trials: $n \in \{1, 2, \ldots\}$

probability of success for j: $\pi_j \in [0, 1]$; $\sum_{j=1}^{k} \pi_j = 1$

$p(\mathbf{y}|n, \boldsymbol{\pi}) = \frac{n!}{y_1! y_2! \ldots y_k!} \pi_1^{y_1} \pi_2^{y_2} \ldots \pi_k^{y_k}$

$E(Y_j) = n\pi_j$

$\text{Var}(Y_j) = n\pi_j(1 - \pi_j)$

`rmultinom(100, size = 5, prob = c(.2,.4,.3,.1))`

What is this?

It's a Prussian soldier getting kicked in the head by a horse.

Also, it's the logo for Stata Press

The Poisson Distribution

- **The story:** count the number of times an (uncommon) event happens
- One parameter: the rate of occurrence, usually called λ
- Represents the number of events occurring in a fixed period of time or in a specific distance, area or volume.
- Can never be negative – so, good for modeling events.
- Assumes that you have lots of trials (or lots of "opportunities" for an event to happen)
- Assumes the probability of a "success" on any particular trial is tiny
- Makes a potentially strong assumption about the mean and variance

- Examples:
    - Number of raindrops that hit a 1" x 1" square on the sidewalk during a rainstorm
    - Number of executive orders a president issues in a week
    - Number Prussian solders who died each year by being kicked in the head by a horse (Bortkiewicz, 1898)

The Poisson Distribution

Histogram of Poisson(5)

$Y \sim \text{Poisson}(\lambda)$

$y = 0, 1, \ldots$

expected number of occurrences, λ is always greater than zero

$p(y|\lambda) = \frac{e^{-\lambda}\lambda^y}{y!}$

$E(Y) = \lambda$

$\text{Var}(Y) = \lambda$

```
rpois(100, lambda = 2)
```

Outline

The Univariate Normal Distribution

- **The story:** any outcome that can take any real number as a value

- Describes data that cluster in a bell curve around the mean.

- It's tough to think of examples of things are are *truly* unbounded, so be careful not to extrapolate your results outside of the range of valid outcomes.

  If we're using a normal distribution to model vote outcomes, don't tell me that you predict a candidate to get 110% of the vote (unless you're studying Pakistan or Vladimir Putin)

- Examples:
  - the heights of male students in our class
  - high school students' SAT scores

The Univariate Normal Distribution

Normal Density

$Y \sim \text{Normal}(\mu, \sigma^2)$

$y \in \mathbb{R}$

mean: $\mu \in \mathbb{R}$

variance: $\sigma^2 > 0$

$p(y|\mu, \sigma^2) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(y-\mu)^2}{2\sigma^2}\right)$

$E(Y) = \mu$

$\text{Var}(Y) = \sigma^2$

```
rnorm(100, mean = 0, sd = 1)
```

# The Uniform Distribution

- **The story:** Any value in the interval you chose is equally probable.

- Two parameters: a and b (or α and β), lower and upper bounds of the range of possible results

- Intuitively easy to understand, but often examples are discrete

- Examples:
  - the degree of longitude you're pointing to if you stop a spinning globe with your finger
  - the number of a person who comes in first in a race (discrete)
  - the lottery tumblers out of which a person draws one ball with a number on it (also discrete)

The Uniform Distribution

Uniform Density

$Y \sim \text{Uniform}(\alpha, \beta)$

$y \in [\alpha, \beta]$

Interval: $[\alpha, \beta]; \ \ \beta > \alpha$

$p(y|\alpha, \beta) = \frac{1}{\beta - \alpha}$

$E(Y) = \frac{\alpha + \beta}{2}$

$\text{Var}(Y) = \frac{(\beta - \alpha)^2}{12}$

```
runif(100, min = -5, max =
10)
```

The Exponential Distribution

- **The story:** how long do you have to wait until an event occurs?

- One parameter: λ, arrival rate of the event

- The distribution assumes that your process is memoryless. The expected time until the event happens is constant, regardless of how much time has passed since the last event.
  $E(y) = E(y|y > 1) = E(y|y > 10)...etc. = 1/\lambda$

- Examples:
  - How long until the bus arrives?
  - Time between bombings in a war-torn country

The Exponential Distribution

**Exponential Density**

$Y \sim \text{Expo}(\lambda)$

$y \in [0, \infty]$

$\lambda > 0$

$p(y|\lambda) = \lambda e^{-\lambda y}$

$E(Y) = \frac{1}{\lambda}$

$\text{Var}(Y) = \frac{1}{\lambda^2}$

```
rexp(100, rate = 3)
```

Outline

1. Replication Paper and other logistics

2. Data Generation Processes and Probability Distributions

3. Discrete Distributions

4. Continuous Distributions

5. **Simulating from Distributions**

6. Distribution Transformations

Coding a density function from scratch

In later problem sets, you're going to have to code likelihood functions in R.

Often you'll be able to use canned functions like `rnorm()` or `rpois()`, but sometimes your likelihood function won't look like a distribution you're familiar with and you'll have to code the density from scratch.

We're going to get practice coding from scratch by programming the PDF of the normal distribution.

To do this we have to first write a function which takes the arguments y, μ, and σ.

Coding the PDF of the normal distribution

Recall that the PDF of a normal distribution is:

$$p(y|\mu, \sigma^2) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(y-\mu)^2}{2\sigma^2}\right)$$

If we want to code this into R, we to set up a new function which takes the arguments y, μ, and σ:

```
normal <- function(y, mu, sigma){ # y must go first!


  exp(-(y - mu) ^ 2 / (2 * sigma^2)) / (sigma * sqrt(2 * pi))
}
```

Coding the PDF of the normal distribution

We know that, by definition, PDFs must integrate to 1 across their support
Let's check that we coded our function correctly by using the

```
integrate() function in R:
 integrate(normal, lower = -1000, upper = 1000,

          mu = 0, #extra arguments needed for your function
          sigma = 1)
 1 with absolute error < 9e-05
```
Notice that we didn't need to integrate from $-\infty$ to ∞ to get the right

answer.

Simulating from a coded PDF

Now we want to simulate from our function and plot its PDF. How could we do this?

To do this we first calculate the density at a bunch of different values of y:

```
values <- seq(-10, 10, .001)
weights <- normal(values,
                  mu = 0,
                  sigma=1)
```

Simulating from a coded PDF

Now we draw samples from our
`values` vector, where the
probability of drawing each value is
weighted by the densities (`weights`)

```
draws <- sample(values,
                size = 10000,
                prob = weights,
                replace = T)
```

We now have a vector of 10000
draws from the PDF we coded. We
use them to plot the PDF

Histogram of simulations from our normal function
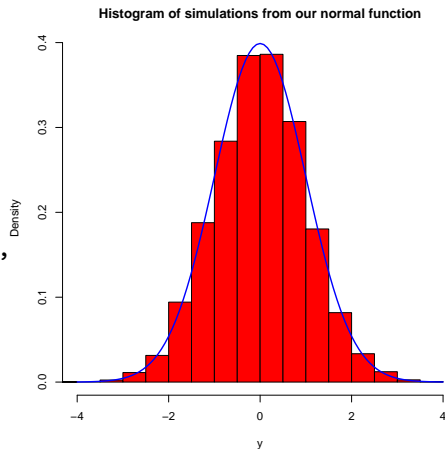
Using simulations to integrate

The draws we took from our function can also be used to integrate
We could use the integrate() function, but for some complicated

likelihoods it might be very slow or not work
```
integrate(normal, lower = -1.96, upper = 1.96,

          mu = 0,
          sigma = 3)
0.9500042 with absolute error < 1e-11
```

Or we can use the draws we took:
```
mean(draws > -1.96 & draws < 1.96)

[1] 0.9506
```

Outline

Transforming Distributions

- $X \sim p(x|\theta)$
- $y = g(x)$
- How is y distributed?
  For example, if $X \sim Exponential(\lambda = 1)$ and $y = log(x)$

$$y \sim ?$$

Transforming Distributions

- It is NOT true that $p(y|\theta) \sim g(p(x|\theta))$. Why?

Transforming Distributions

The Rule

- $X \sim p_x(x|\theta)$
- $y = g(x)$

$$p_y(y) = p_x(g^{-1}(y)) \left| \frac{dg^{-1}}{dy} \right|$$

- What is $g^{-1}(y)$? The inverse of y=g(x).
- What is $\left| \frac{dg^{-1}}{dy} \right|$? The Jacobian.

Transforming Distributions – the log-Normal Example

For example,

- $X \sim Normal(x|\mu = 0, \sigma = 1)$
- $y = g(x) = e^x$
- what is $g^{-1}(y)$?

$$g^{-1}(y) = x = log(y)$$

- What is $\frac{dg^{-1}}{dy}$?

$$\frac{d(log(y))}{dy} = \frac{1}{y}$$

Transforming Distributions – the log-Normal Example

- Put it all together

$$p_y(y) \;=\; p_x\left(log(y)\right)\left|\frac{1}{y}\right|$$

- Notice we don't need the absolute value because $y > 0$.

$$p_y(y) \;=\; \frac{1}{\sqrt{2\pi}}e^{-\frac{1}{2}(log(y))^2}\frac{1}{y}$$

- $Y \sim$ log-Normal$(0, 1)$
- Challenge: derive the chi-squared distribution.
  $X \sim N(\mu, \sigma^2)$
  $Y = X^2$